

GRHydro: A new open source general-relativistic magnetohydrodynamics code for the Einstein Toolkit

Philipp Mösta¹, Bruno C. Mundim^{2,3}, Joshua A. Faber³,
Roland Haas^{1,4}, Scott C. Noble³, Tanja Bode^{8,4}, Frank Löffler⁵,
Christian D. Ott^{1,5}, Christian Reisswig¹, Erik Schnetter^{6,7,5}

¹ TAPIR, California Institute of Technology, Pasadena, CA 91125, USA

² Max-Planck-Institut für Gravitationsphysik, Albert-Einstein-Institut, D-14476
Golm, Germany

³ Center for Computational Relativity and Gravitation and School of Mathematical
Sciences, Rochester Institute of Technology, Rochester, NY 14623, USA

⁴ Center for Relativistic Astrophysics, School of Physics, Georgia Institute of
Technology, Atlanta, GA 30332, USA

⁵ Center for Computation & Technology, Louisiana State University, Baton Rouge,
LA 70803, USA

⁶ Perimeter Institute for Theoretical Physics, Waterloo, ON N2L 2Y5, Canada

⁷ Department of Physics, University of Guelph, Guelph, ON N1G 2W1, Canada

⁸ Theoretische Astrophysik, Institut für Astronomie und Astrophysik, Universität
Tübingen, 72076 Tübingen, Germany

E-mail: pmoesta@tapir.caltech.edu

Abstract. We present the new general-relativistic magnetohydrodynamics (GRMHD) capabilities of the Einstein Toolkit, an open-source community-driven numerical relativity and computational relativistic astrophysics code. The GRMHD extension of the Toolkit builds upon previous releases and implements the evolution of relativistic magnetised fluids in the ideal MHD limit in fully dynamical spacetimes using the same shock-capturing techniques previously applied to hydrodynamical evolution. In order to maintain the divergence-free character of the magnetic field, the code implements both hyperbolic divergence cleaning and constrained transport schemes. We present test results for a number of MHD tests in Minkowski and curved spacetimes. Minkowski tests include aligned and oblique planar shocks, cylindrical explosions, magnetic rotors, Alfvén waves and advected loops, as well as a set of tests designed to study the response of the divergence cleaning scheme to numerically generated monopoles. We study the code’s performance in curved spacetimes with spherical accretion onto a black hole on a fixed background spacetime and in fully dynamical spacetimes by evolutions of a magnetised polytropic neutron star and of the collapse of a magnetised stellar core. Our results agree well with exact solutions where these are available and we demonstrate convergence. All code and input files used to generate the results are available on <http://einstein toolkit.org>. This makes our work fully reproducible and provides new users with an introduction to applications of the code.

1. Introduction

The recent years have seen rapid developments in the field of numerical relativity. Beginning with the first fully general-relativistic (GR) simulations of binary neutron star (NS) mergers by Shibata and Uryu in 1999 [1], fully dynamical general-relativistic hydrodynamics (GRHD) has been explored by a growing number of research groups. A major step forward occurred 2005, when three independent groups developed two different techniques, the generalised harmonic gauge formalism [2] and the so-called “moving puncture” method [3, 4], to evolve vacuum spacetimes containing black holes (BH) without encountering numerical instabilities that mark the end of a numerical simulation. Since then, these evolution schemes have been incorporated in hydrodynamic simulations and now one or the other is used in essentially every full GR calculation of merging binaries, as well as other astrophysical phenomena involving dynamical spacetimes, such as stellar collapse and BH formation (e., [5]).

A more recent advance concerns the incorporation of magnetic fields, particularly in the ideal magnetohydrodynamics (MHD) limit, into dynamical simulations. Building upon work that originally focused on astrophysical systems in either Newtonian gravity or fixed GR backgrounds, as is appropriate for studies of accretion disks (see [6–8] for reviews of the topic), MHD evolution techniques have been incorporated into dynamical GR codes and used to study the evolution of NS-NS (e.g., [9–12]) and BH-NS mergers (e.g., [13, 14]), self-gravitating tori around black holes [15, 16], neutron star collapse (e.g., [17, 18]), stellar core collapse [19], and the evolution of magnetised plasma around merging binary BHs (e.g., [20, 21]). Given their rapidly maturing capabilities and widespread use, many numerical relativity codes have been made public, either in total or in part. The largest community-based effort to do so is the Einstein Toolkit consortium [22, 23], which has as one of its goals to provide a free, publicly available, community-driven general-relativistic MHD (GRMHD) code capable of performing simulations that include realistic treatments of matter, electromagnetic fields, and gravity. The code is built upon several open-source components that are widely used throughout the numerical relativity community, including the **Cactus** computational infrastructure [24–26], the **Carp** adaptive mesh refinement (AMR) code [27–29], the **McLachlan** GR evolution code [30, 31], and the GRHD code **GRHydro** [22], which has been developed starting from a public variant of the **Whisky** code [32–35]. Many of the details of the Einstein Toolkit may be found in [22], which describes the routines used to provide the supporting computational infrastructure for grid setup and parallelization, constructing initial data, evolving dynamical GRHD configurations, and analysing simulation outputs. An extension of the Einstein Toolkit to general multi-patch grids with Cartesian and curvilinear geometry is discussed in [36].

In this paper, we describe the newest component of the Einstein Toolkit: a MHD evolution scheme incorporated into the **GRHydro** module. Using what is known as the *Valencia* formulation [37–39] of the GRMHD equations, the new code evolves magnetic fields in fully dynamical GR spacetimes under the assumptions of ideal MHD, i.e., the

resistivity is taken to be zero and electric fields vanish in the comoving frame of the fluid. While the evolution equations themselves are easily cast into the same flux-conservative form as the other hydrodynamics equations, the primary challenge for evolving magnetic fields is numerically maintaining the divergence-free constraint of the magnetic field. For this, we have implemented a variant of the constrained transport approach [40] and a hyperbolic “divergence cleaning” technique, similar to that discussed in [41, 42].

The `GRHydro` module of the Einstein Toolkit is a new, independently written code, and—in particular—the development is completely independent of the `WhiskyMHD` code [43], since the `GRHydro` development started from the GRHD (non-MHD) version of the `Whisky` code. Given its close interaction with the hydrodynamic evolution, the MHD routines are contained within the `GRHydro` package (or “thorn”, in the language of `Cactus`-based routines). It is developed in a public repository, along with other components of the Einstein Toolkit, and it is publicly available under the same open-source licensing terms as other components in the Toolkit. Given the public nature of the project, the code release includes the subroutines themselves, complete documentation for their use, and parameter files needed to reproduce the tests shown here. Code and parameter files are available on the Einstein Toolkit web page, <http://einsteintoolkit.org>.

In order to validate our GRMHD implementation, we include a number of flat-space tests whose solutions are either known exactly or approximately, and which have frequently been used as a testbed for other GRMHD codes. These include planar and cylindrical shocktubes, a rotating bar threaded by a magnetic field (the “magnetic rotor”), propagating Alfvén waves, and the advection of a flux loop. We test the code’s ability to handle static curved spacetime by studying Bondi accretion onto a Schwarzschild BH and perform GRMHD simulations with dynamical spacetime evolution for a magnetised TOV star and for the collapse of a rotating magnetised stellar core to a rotating neutron star.

This paper is structured as follows: in section 2, we describe the Einstein Toolkit and briefly discuss other relativistic MHD codes. In section 3, we describe the Valencia formulation of the GRMHD equations, and in section 4 the numerical techniques used in `GRHydro`. In section 5, we describe the tests we have carried out with the new code. Finally, in section 6, we summarise and discuss future directions of the Einstein Toolkit.

2. Numerical Relativistic Hydrodynamics and GRMHD codes

The `GRHydro` code is, to the best of our knowledge, the first publicly released 3D GRMHD code capable of evolving configurations in fully dynamical spacetimes. Still, its development makes use of code and techniques from a number of other public codes, which we summarise briefly here. The most obvious of these is the Einstein Toolkit, within which its development has taken place. As we discuss below, the Toolkit itself is composed of tens of different modules developed by a diverse set of authors whose number is approaching 100. The MHD techniques, while independently implemented in

nearly all cases, rely heavily on both the numerical techniques found within the **GRHydro** package [22], and thus the original **Whisky** code [32–35], as well as the publicly available **HARM** code [44].

2.1. The Einstein Toolkit

The Einstein Toolkit includes components to aid in performing relativistic astrophysical simulations that range from physics modelling (initial conditions, evolution, analysis) to infrastructure modules (grid setup, parallelization, I/O) and include related tools (workflow management, file converters, etc.). The overall goal is to provide a set of well-documented, well-tested, state-of-the-art components for these tasks, while allowing users to replace these components with their own and/or add additional ones.

Many components of the Einstein Toolkit use the **Cactus** Computational Toolkit [24–26], a software framework for high-performance computing. **Cactus** simplifies designing codes in a modular (“component-based”) manner, and many existing **Cactus** modules provide infrastructure facilities or basic numerical algorithms such as coordinates, boundary conditions, interpolators, reduction operators, or efficient I/O in different data formats.

The utilities contained in Einstein Toolkit, e.g., help manage components [45, 46], build code and submit simulations on supercomputers [47, 48], or provide remote debuggers [49] and post-processing and visualisation interfaces for **VisIt** [50].

Adaptive mesh refinement (AMR) and multi-block methods are implemented by **Carpenter** [27–29] which also provides MPI parallelization. **Carpenter** supports Berger-Oliger style (“block-structured”) adaptive mesh refinement [51] with sub-cycling in time as well as certain additional features commonly used in numerical relativity (see [27] for details). **Carpenter** provides also the respective prolongation (interpolation) and restriction (projection) operators to move data between coarse and fine grids. **Carpenter** has been demonstrated to scale efficiently up to several thousand cores [22].

Carpenter supports both vertex-centred and cell-centred AMR. In vertex-centred AMR, each coarse grid point (cell Center) coincided with a fine grid point (cell Center), which simplifies certain operations (e.g. restriction, and also visualisation). In this paper, we present results obtained with vertex-centred AMR only. In cell-centred AMR, coarse grid *cell faces* are aligned with fine grid faces, allowing in particular exact conservation across AMR interfaces. For a more detailed discussion of cell-centred AMR in hydrodynamics simulations with the Einstein Toolkit we refer the reader to [36].

The evolution of the spacetime metric in the Einstein Toolkit is handled by the **McLachlan** package [30, 31]. This code is autogenerated by Mathematica using the **Kranc** package [52–54], implementing the Einstein equations via a 3 + 1-dimensional split using the BSSN formalism [55–59]. The BSSN equations are finite-differenced at a user-specified order of accuracy, and coupling to hydrodynamic variables is included via the stress-energy tensor. The time integration and coupling with curvature are carried out with the Method of Lines (MoL) [60], implemented in the **MoL** package.

Hydrodynamic evolution techniques are provided in the Einstein Toolkit by the **GRHydro** package, a code derived from the public **Whisky** GRHD code [32–35]. The code is designed to be modular, interacting with the vacuum metric evolution only by contributions to the stress-energy tensor and by the local values of the metric components and extrinsic curvature. It uses a high-resolution shock capturing finite-volume scheme to evolve hydrodynamic quantities, with several different methods for reconstruction of states on cell interfaces and Riemann solvers. It also assumes an atmosphere to handle low-density and vacuum regions. In particular, a density floor prevents numerical errors from developing near the edge of matter configurations. Boundaries and symmetries are handled by registering hydrodynamic variables with the appropriate **Cactus** routines. Passive “tracers” or local scalars that are advected with the fluid flow may be used, though their behaviour is generally unaffected by the presence or absence of magnetic fields.

In adding MHD to the pre-existing hydrodynamics code, our design philosophy has been to add the new functionality in such a way that the original code would run without alteration should MHD not be required. For nearly all of the more complicated routines in the code, this involved creating parallel MHD and non-MHD routines, while for some of the more basic routines we simply branch between different code sections depending on whether MHD is required or not. While this requires care when the code is updated, since some changes may need to be implemented twice, it does help to insure that users performing non-MHD simulations will be protected against possible errors introduced in the more frequently changing MHD routines.

2.2. Other relativistic MHD codes

In extending the Einstein Toolkit to include MHD functionality, we incorporated techniques that have previously appeared in the literature. Of particular importance is the **HARM** code [44, 61, 62], a free, publicly available GRMHD code that can operate on fixed background metrics, particularly those describing spherically symmetric or rotating black holes. Our routines for converting conservative variables back into primitive ones are adapted directly from **HARM** ([63]; see section 4.4 below), as is the approximate technique used to calculate wave speeds for the Riemann solver ([44]; see section 4.3 below).

Since many groups have introduced relativistic MHD codes, there are a number of standard tests that may be used to gauge the performance of a particular code. Many of these are compiled in the descriptions of the **HARM** code mentioned above, as well as papers describing the **Athena** MHD code [64, 65], the **Echo** GRMHD code [66], the Tokyo/Kyoto group’s GRMHD code [67], the UIUC GRMHD code [68, 69], the **WhiskyMHD** code [43], and the LSU GRMHD code [70]. We have implemented several of these tests, covering all aspects of our code, as we describe in section 5 below.

3. The Valencia formulation of ideal MHD

While **GRHydro** does not technically assume a particular evolution scheme to evolve the field equations for the GR metric, the most widely used choice for the ET code is the BSSN formalism [56, 57], the particular implementation of which is provided by the **McLachlan** code [30, 31]. We do assume that the metric is known in the ADM (Arnowitt-Deser-Misner) form [71], in which we have

$$ds^2 = g_{\mu\nu}dx^\mu dx^\nu \equiv (-\alpha^2 + \beta_i\beta^i)dt^2 + 2\beta_i dt dx^i + \gamma_{ij}dx^i dx^j, \quad (1)$$

with $g_{\mu\nu}$, α , β^i , and γ_{ij} being the spacetime 4-metric, lapse function, shift vector, and spatial 3-metric, respectively. Note that we are assuming spacelike signature, so the Minkowski metric in flat space reads $\eta_{\mu\nu} = \text{diag}(-1, 1, 1, 1)$. Roman indices are used for 3-quantities and 4-quantities are indexed with Greek characters. We work in units of $c = G = M_\odot = 1$ unless explicitly stated otherwise.

GRHydro employs the *ideal MHD* approximation – fluids have infinite conductivity and there is no charge separation. Thus, electric fields $E_\nu = u_\mu F^{\mu\nu}$ in the rest frame of the fluid vanishes and ideal MHD corresponds to imposing the following conditions:

$$u_\mu F^{\mu\nu} = 0. \quad (2)$$

Note that throughout this paper we rescale the magnitude of the relativistic Faraday tensor $F^{\mu\nu}$ and its dual ${}^*F^{\mu\nu} \equiv \frac{1}{2}\epsilon^{\mu\nu\kappa\lambda}F_{\kappa\lambda}$, as well as the magnetic and electric fields, by a factor of $1/\sqrt{4\pi}$ to eliminate the need to include the permittivity and permeability of free space in cgs-Gaussian units.

Magnetic fields enter into the equations of hydrodynamics by their contributions to the stress-energy tensor and through the solutions of Maxwell's equations. The hydrodynamic and electromagnetic contributions to the stress-energy tensor are given, respectively, by

$$T_H^{\mu\nu} = \rho h u^\mu u^\nu + P g^{\mu\nu} = (\rho + \rho\epsilon + P) u^\mu u^\nu + P g^{\mu\nu}, \quad (3)$$

and

$$T_{\text{EM}}^{\mu\nu} = F^{\mu\lambda} F^\nu{}_\lambda - \frac{1}{4} g^{\mu\nu} F^{\lambda\kappa} F_{\lambda\kappa} = b^2 u^\mu u^\nu - b^\mu b^\nu + \frac{b^2}{2} g^{\mu\nu}, \quad (4)$$

where ρ , ϵ , P , u^μ , and $h \equiv 1 + \epsilon + P/\rho$ are the fluid rest mass density, specific internal energy, gas pressure, 4-velocity, and specific enthalpy, respectively, and b^μ is the magnetic 4-vector (the projected component of the Maxwell tensor parallel to the 4-velocity of the fluid):

$$b^\mu = u_\nu {}^*F^{\mu\nu}. \quad (5)$$

Note that $b^2 = b^\mu b_\mu = 2P_m$, where P_m is the magnetic pressure. Combined, the stress-energy tensor takes the form:

$$\begin{aligned} T^{\mu\nu} &= (\rho + \rho\epsilon + P + b^2) u^\mu u^\nu + \left(P + \frac{b^2}{2}\right) g^{\mu\nu} - b^\mu b^\nu \\ &\equiv \rho h^* u^\mu u^\nu + P^* g^{\mu\nu} - b^\mu b^\nu, \end{aligned} \quad (6)$$

where we define the magnetically modified pressure and enthalpy, $P^* = P + P_m = P + b^2/2$ and $h^* \equiv 1 + \epsilon + (P + b^2)/\rho$, respectively.

The spatial magnetic field (living on the spacelike 3-hypersurfaces), is defined as the Eulerian component of the Maxwell tensor

$$B^i = n_\mu {}^*F^{i\mu} = -\alpha {}^*F^{i0} , \quad (7)$$

where $n_\mu = [-\alpha, 0, 0, 0]$ is the normal vector to the hypersurface.

The equations of ideal GRMHD evolved by **GRHydro** are derived from the local GR conservation laws of mass and energy-momentum,

$$\nabla_\mu J^\mu = 0 , \quad \nabla_\mu T^{\mu\nu} = 0 , \quad (8)$$

where ∇_μ denotes the covariant derivative with respect to the 4-metric and $J^\mu = \rho u^\mu$ is the mass current, and from Maxwell's equations,

$$\nabla_\nu {}^*F^{\mu\nu} = 0 . \quad (9)$$

The **GRHydro** scheme is written in a first-order hyperbolic flux-conservative evolution system for the conserved variables D , S^i , τ , and \mathcal{B}^i , defined in terms of the primitive variables ρ, ϵ, v^i , and B^i such that

$$D = \sqrt{\gamma} \rho W , \quad (10)$$

$$S_j = \sqrt{\gamma} (\rho h^* W^2 v_j - \alpha b^0 b_j) , \quad (11)$$

$$\tau = \sqrt{\gamma} (\rho h^* W^2 - P^* - (\alpha b^0)^2) - D , \quad (12)$$

$$\mathcal{B}^k = \sqrt{\gamma} B^k , \quad (13)$$

where γ is the determinant of γ_{ij} . We choose a definition of the 3-velocity v^i that corresponds to the velocity seen by an Eulerian observer at rest in the current spatial 3-hypersurface [72],

$$v^i = \frac{u^i}{W} + \frac{\beta^i}{\alpha} , \quad (14)$$

and $W \equiv (1 - v^i v_i)^{-1/2}$ is the Lorentz factor. Note that v^i , B^i , S^i , and β^i are 3-vectors, and their indices are raised and lowered with the 3-metric, e.g., $v_i \equiv \gamma_{ij} v^j$.

The evolution scheme used in **GRHydro** is often referred to as the Valencia formulation [6, 37–39]. Our notation here most closely follows that found in [73]. The evolution system for the conserved variables, representing (8) and the spatial components of (9), is

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}^i}{\partial x^i} = \mathbf{S} , \quad (15)$$

with

$$\begin{aligned} \mathbf{U} &= [D, S_j, \tau, \mathcal{B}^k] , \\ \mathbf{F}^i &= \alpha \times \begin{bmatrix} D \tilde{v}^i \\ S_j \tilde{v}^i + \sqrt{\gamma} P^* \delta_j^i - b_j \mathcal{B}^i / W \\ \tau \tilde{v}^i + \sqrt{\gamma} P^* v^i - \alpha b^0 \mathcal{B}^i / W \\ \mathcal{B}^k \tilde{v}^i - \mathcal{B}^i \tilde{v}^k \end{bmatrix} , \end{aligned} \quad (16)$$

$$\mathbf{S} = \alpha\sqrt{\gamma} \times \begin{bmatrix} 0, \\ T^{\mu\nu} \left(\frac{\partial g_{\nu j}}{\partial x^\mu} - \Gamma_{\mu\nu}^\lambda g_{\lambda j} \right) \\ \alpha \left(T^{\mu 0} \frac{\partial \ln \alpha}{\partial x^\mu} - T^{\mu\nu} \Gamma_{\mu\nu}^0 \right) \\ \vec{0} \end{bmatrix}. \quad (17)$$

Here, $\tilde{v}^i = v^i - \beta^i/\alpha$ and $\Gamma_{\mu\nu}^\lambda$ are the 4-Christoffel symbols.

The time component of (9) yields the condition that the magnetic field is divergence-free, the “no-monopoles” constraint:

$$\nabla \cdot B \equiv \frac{1}{\sqrt{\gamma}} \partial_i (\sqrt{\gamma} B^i) = 0, \quad (18)$$

which also implies

$$\partial_i \mathcal{B}^i = 0. \quad (19)$$

In practice, we implement two different methods to actively enforce this constraint. In the “divergence cleaning” technique we include the ability to modify the magnetic field evolution by introducing a new field variable that dissipates away numerical divergences, which we discuss in detail in section 4.5.1 below. An alternative method, commonly called “constrained transport”, instead carefully constructs a numerical method so that the constraint (19) is satisfied to round-off error at the discrete level. We discuss this method in section 4.5.2.

4. Numerical Methods

GRHydro’s GRMHD code uses the same infrastructure and backend routines as its pure general relativistic hydrodynamics variant [22]. In the following, we focus on the discussion of the numerical methods used in the extension to GRMHD.

GRHydro’s GRMHD code implements reconstruction of fluid and magnetic field variables to cell interfaces for all of the methods present in the original GRHydro code: TVD (total variation diminishing) (e.g., [74]), PPM (piecewise parabolic method) [75], and ENO (essentially non-oscillatory) [76, 77]. In addition, we added enhanced PPM (ePPM, as described in [36, 78]), WENO5 (5th order weighted-ENO) [79], and MP5 (5th order monotonicity preserving) [80] both in GRMHD and pure GR hydrodynamics simulations. We discuss the different reconstruction methods in section 4.2. The code computes the solution of the local Riemann problems at cell interfaces using the HLLE (Harten-Lax-van Leer-Einfeldt) [81, 82] approximate Riemann solver discussed in section 4.3. We implement the conversion of conserved to primitive variables for arbitrary equations of state (EOS), including polytropic, Γ -law, hybrid polytropic/ Γ -law, and microphysical finite-temperature EOS. We summarise the new methods for GRMHD in section 4.4. An important aspect of any (GR)MHD scheme is the numerical method used to preserve the divergence free constraint. GRHydro implements both the hyperbolic divergence cleaning method (e.g., [41, 42]) and a variant of the constrained transport method (e.g., [40, 43, 83]), which are both discussed in detail in section 4.5.

4.1. Evaluation of magnetic field expressions

The MHD code stores the values of both the primitive B-field vector, B^i (`Bvec` in the code), and the evolved conservative field \mathcal{B}^i (`Bcons` in the code) in the frame of the Eulerian observers. For analysis purposes we include options to compute the magnetic field in fluid's rest frame

$$b^0 = \frac{WB^k v_k}{\alpha}, \quad (20)$$

$$b^i = \frac{B^i}{W} + W(B^k v_k) \left(v^i - \frac{\beta^i}{\alpha} \right), \quad (21)$$

$$b^2 = \frac{B^i B_i}{W^2} + (B^i v_i)^2. \quad (22)$$

$$(23)$$

4.2. Reconstruction

In a finite-volume scheme, one evaluates fluxes at cell faces by solving Riemann problems involving potentially discontinuous hydrodynamic states on either side of the interface. To construct these Riemann problems, one must first obtain the fluid state on the left and right sides of the interface. The fluid state is known within cells in the form of cell-averages of the hydrodynamical variables. The *reconstruction* step interpolates the fluid state from cell averaged values to values at cell interfaces without introducing oscillations at shocks and other discontinuities. It is possible to reconstruct either primitive or conserved fluid variables, however using the former makes it much easier to guarantee physically valid results for which the pressure is positive and the fluid velocities sub-luminal.

By assuming an orthogonal set of coordinates, it is possible to reconstruct each coordinate direction independently. Thus, the fluid reconstruction reduces to a one-dimensional problem.

We define $U_{i+1/2}^L$ to be the value of an element of our conservative variable state vector \mathbf{U} on the left side of the face between $U_i \equiv U(x_i, y, z)$ and $U_{i+1} \equiv U(x_{i+1}, y, z)$, where x_i is the i^{th} point in the x -direction, and $U_{i+1/2}^R$ the value on the right side of the same face (coincident in space, but with a potentially different value). These quantities are computed directly from the primitive values on the face.

4.2.1. TVD reconstruction. For total variation diminishing (TVD) methods, we let

$$U_{i+1/2}^L = U_i + \frac{F(U_i)\Delta x}{2}; \quad U_{i+1/2}^R = U_{i+1} - \frac{F(U_{i+1})\Delta x}{2} \quad (24)$$

where $F(U_i)$ is a slope-limited gradient function, typically determined by the values of $U_{i+1} - U_i$ and $U_i - U_{i-1}$, with a variety of different forms of the slope limiter available. In practice, all try to accomplish the same task of preserving monotonicity and removing the possibility of spuriously creating local extrema. `GRHydro` includes minmod, superbee [84], and monotonized central [85] limiters. TVD methods are second-order

accurate in regions of smooth, monotonic flows. At extrema and shocks, they reduce to first order.

4.2.2. PPM reconstruction. The original piecewise parabolic method (oPPM, [75]) uses quadratic functions to represent cell-averages from which new states at cell interfaces are constructed. The fluid state is interpolated using a fourth-order polynomial. A number of subsequent limiter steps constrain parabolic profiles and preserve monotonicity so that no new extrema can form. The version implemented in **GRHydro** includes the steepening and flattening routines described in the original PPM papers, with a simplified flattening procedure that allows for fewer ghost points [32, 36]. The original PPM always reduces to first order near local extrema and shocks. The enhanced PPM (ePPM) maintains high-order at local extrema that are smooth [36, 78].

4.2.3. ENO reconstruction. Essentially non-oscillatory (ENO) methods use a divided differences approach to achieve high-order accuracy via polynomial interpolation [76, 77] without reducing the order near extrema or shocks. In the third-order case, two interpolation polynomials with different stencil points are used. Based on the smoothness of the field, one of the interpolation polynomials is selected.

4.2.4. WENO reconstruction Weighted essentially non-oscillatory (WENO) reconstruction [79] is an improved algorithm based on the ENO approach. The drawback of ENO methods is that the order of accuracy is not maximal for the available number of stencil points. Furthermore, since a single stencil is selected, ENO reconstruction is not continuous. In addition, the large number of required *if statements* make ENO methods unnecessary slow. WENO reconstruction, on the other hand, attempts to overcome all these drawbacks. The essential idea is to combine all possible ENO reconstruction stencils by assigning a weight to each stencil. The weight is determined by the local smoothness of the flow. When all weights are non-zero, the full set of stencil points is used, and the maximum allowable order of accuracy for a given number of stencil points is achieved. When the flow becomes less smooth, some weights are suppressed, and a particular lower order interpolation stencil dominates the reconstruction. Using the same number of stencil points as third-order ENO, WENO is fifth-order accurate when the flow is smooth, and reduces to third order near shocks and discontinuities.

We implement fifth-order WENO reconstruction based on an improved version presented in [86]. We introduce three interpolation polynomials that approximate U_{i+1}^L from cell-averages of a given quantity U_i :

$$U_{i+1/2}^{L,1} = \frac{3}{8}U_{i-2} - \frac{5}{4}U_{i-1} + \frac{15}{8}U_i, \quad (25)$$

$$U_{i+1/2}^{L,2} = -\frac{1}{8}U_{i-1} + \frac{3}{4}U_i - \frac{3}{8}U_{i+1}, \quad (26)$$

$$U_{i+1/2}^{L,3} = \frac{3}{8}U_i + \frac{3}{4}U_{i+1} - \frac{1}{8}U_{i+2}. \quad (27)$$

Each of the three polynomials yields a third-order accurate approximation of \mathbf{U} at the cell-interface. By introducing a convex linear combination of the three interpolation polynomials,

$$U_{i+1/2}^L = w^1 U_{i+1/2}^{L,1} + w^2 U_{i+1/2}^{L,2} + w^3 U_{i+1/2}^{L,3} , \quad (28)$$

where the weights w^i satisfy $\sum_i w^i = 1$, it is possible to obtain a fifth-order interpolation polynomial that spans all five stencil points $\{U_{i-2}, \dots, U_{i+2}\}$. The weights w^i are computed using so-called smoothness indicators β^i . In the original WENO algorithm [79], they are given by

$$\beta^1 = \frac{1}{3}(4u_{i-2}^2 - 19u_{i-2}u_{i-1} + 25u_{i-1}^2 + 11u_{i-2}u_i - 31u_{i-1}u_i + 10u_i^2) , \quad (29)$$

$$\beta^2 = \frac{1}{3}(4u_{i-1}^2 - 13u_{i-1}u_i + 13u_i^2 + 5u_{i-1}u_{i+1} - 13u_iu_{i+1} + 4u_{i+1}^2) , \quad (30)$$

$$\beta^3 = \frac{1}{3}(10u_i^2 - 31u_iu_{i+1} + 25u_{i+1}^2 + 11u_iu_{i+2} - 19u_{i+1}u_{i+2} + 4u_{i+2}^2) . \quad (31)$$

Note that, in this section only, β^i refers to the smoothness indicators (29) and not to the shift vector (1). The weights are then obtained via

$$w^i = \frac{\bar{w}^i}{\bar{w}^1 + \bar{w}^2 + \bar{w}^3} , \quad \text{with} \quad \bar{w}^i = \frac{\gamma^i}{(\epsilon + \beta^i)^2} , \quad (32)$$

where $\gamma^i = \{1/16, 5/8, 5/16\}$, and ϵ is a small constant to avoid division by zero. Unfortunately, the choice of ϵ is scale dependent, and choosing a fixed number is inappropriate for cases with large variation in scales. The improvement made by [86] overcomes this problem by modifying the smoothness indicators (29)-(31). Following [86], we compute *modified* smoothness indicators via

$$\bar{\beta}^i = \beta^i + \epsilon |U^2| + \delta , \quad (33)$$

where $|U^2|$ is the sum of the U_j^2 in the i -th stencil, and δ is the smallest number that the chosen floating point variable type can hold. Now, the smoothness indicators depend on the scale of the reconstructed field. In our case, we set

$$\bar{\beta}^i = \beta^i + \epsilon(|U^2| + 1) , \quad (34)$$

with $\epsilon = 10^{-26}$ and use these values instead of β^i in (29)-(31).

4.2.5. MP5 reconstruction Monotonicity-preserving fifth-order (MP5) reconstruction is based on a geometric approach to maintain high-order at maxima that are smooth. In contrast to WENO or ENO, MP5 makes use of limiters to a high-order reconstruction polynomial similar to PPM to avoid oscillations at shocks and discontinuities. The key advantages of MP5 reconstruction are that it preserves monotonicity and accuracy, and is fast. In practice, it compares favourable to enhanced PPM and WENO in terms of accuracy (see Appendix B). MP5 reconstruction is carried out in two steps. In a first step, a fifth-order polynomial is used to interpolate cell-averages U_i on cell interfaces $U_{i+1/2}^L$:

$$U_{i+1/2}^L = (2U_{i-2} - 13U_{i-1} + 47U_i + 27U_{i+1} - 3U_{i+2})/60 . \quad (35)$$

In a second step, the interpolated value $U_{i+1/2}^L$ is limited. Whether a limiter is applied is determined via the following condition. First we compute

$$U^{MP} = U_i + \text{minmod}(U_{i+1} - U_i, \tilde{\alpha}(U_i - U_{i-1})) , \quad (36)$$

where $\tilde{\alpha}$ is a constant, which we set to 4.0, and where

$$\text{minmod}(x, y) = \frac{1}{2} (\text{sign}(x) + \text{sign}(y)) \min(|x|, |y|) . \quad (37)$$

A limiter is *not* applied when

$$(U_{i+1/2}^L - U_i)(U_{i+1/2}^L - U^{MP}) \leq \epsilon |U| , \quad (38)$$

where ϵ is a small constant, and $|U|$ is the L_2 norm of U_i over the stencil points $\{U_{i-2}, \dots, U_{i+2}\}$. Note that the norm factor is not present in the original algorithm. We have added this to take into account the scale of the reconstructed field. Following [80], we set $\epsilon = 10^{-10}$, though in some cases, it may be necessary to set $\epsilon = 0$ to avoid oscillations at strong shocks or contact discontinuities such as the surface of a neutron star.

In case (38) does not hold, the following limiter algorithm is applied. First, we compute the second derivatives

$$D_i^- = U_{i-2} - 2U_{i-1} + U_i , \quad (39)$$

$$D_i^0 = U_{i-1} - 2U_i + U_{i+1} , \quad (40)$$

$$D_i^+ = U_i - 2U_{i+1} + U_{i+2} . \quad (41)$$

Next, we compute

$$D_{i+1/2}^{M4} = \text{minmod}(4D_i^0 - D_i^+, 4D_i^+ - D_i^0, D_i^0, D_i^+) , \quad (42)$$

$$D_{i-1/2}^{M4} = \text{minmod}(4D_i^0 - D_i^-, 4D_i^- - D_i^0, D_i^0, D_i^-) , \quad (43)$$

where the four-argument minmod function is given by

$$\begin{aligned} \text{minmod}(w, x, y, z) = \frac{1}{8} & (\text{sign}(w) + \text{sign}(x)) \times \\ & |(\text{sign}(w) + \text{sign}(y))(\text{sign}(w) + \text{sign}(z))| \times \\ & \min(|w|, |x|, |y|, |z|) . \end{aligned} \quad (44)$$

We then compute

$$U^{UL} = U_i + \alpha(U_i - U_{i+1}) , \quad (45)$$

$$U^{AV} = \frac{1}{2}(U_i + U_{i+1}) , \quad (46)$$

$$U^{MD} = U^{AV} - \frac{1}{2}D_{i+1/2}^{M4} , \quad (47)$$

$$U^{LC} = U_i + \frac{1}{2}(U_i - U_{i-1}) + \frac{4}{3}D_{i-1/2}^{M4} . \quad (48)$$

Using these expressions, we compute

$$U_{\min} = \max(\min(U_i, U_{i+1}, U^{MD}), \min(U_i, U^{UL}, U^{LC})) , \quad (49)$$

$$U_{\max} = \min(\max(U_i, U_{i+1}, U^{MD}), \max(U_i, U^{UL}, U^{LC})) . \quad (50)$$

$$(51)$$

Finally, a new limited value for the cell interface $U_{i+1/2}^L$ is obtained via

$$U_{i+1/2}^{L,\text{limited}} = U_{i+1/2}^L + \text{minmod}(U_{\min} - U_{i+1/2}^L, U_{\max} - U_{i+1/2}^L) . \quad (52)$$

To obtain the reconstructed value at the right interface $U_{i-1/2}^R$, the values at U_{i-2}, \dots, U_{i+2} are replaced by the values U_{i+2}, \dots, U_{i-2} .

For a detailed description and derivation of the MP5 reconstruction algorithm, we refer the reader to the original paper [80].

4.3. Riemann Solver

We implement the Harten-Lax-van Leer-Einfeldt (HLLE) approximate solver [82, 87]. The more accurate Roe and Marquina solvers require determining the eigenvalues characterising the linearizing hydrodynamic evolution scheme, which is extremely resource intensive while not providing a decisive advantage in accuracy. In contrast, HLLE uses a two-wave approximation to compute the update terms across the discontinuity at the cell interface. With ξ_- and ξ_+ the most negative and most positive wave speed eigenvalues present on either side of the interface (including magnetic field modes but not those associated with divergence cleaning subsystem, as discussed in section 4.5.1), the solution state vector \mathbf{U} is assumed to take the form

$$\mathbf{U} = \begin{cases} \mathbf{U}^L & \text{if } 0 < \xi_- , \\ \mathbf{U}_* & \text{if } \xi_- < 0 < \xi_+ , \\ \mathbf{U}^R & \text{if } 0 > \xi_+ , \end{cases} \quad (53)$$

with the intermediate state \mathbf{U}_* given by

$$\mathbf{U}_* = \frac{\xi_+ \mathbf{U}^R - \xi_- \mathbf{U}^L - \mathbf{F}(\mathbf{U}^R) + \mathbf{F}(\mathbf{U}^L)}{\xi_+ - \xi_-} . \quad (54)$$

The numerical flux along the interface takes the form

$$\mathbf{F}(\mathbf{U}) = \frac{\hat{\xi}_+ \mathbf{F}(\mathbf{U}^L) - \hat{\xi}_- \mathbf{F}(\mathbf{U}^R) + \hat{\xi}_+ \hat{\xi}_- (\mathbf{U}^R - \mathbf{U}^L)}{\hat{\xi}_+ - \hat{\xi}_-} , \quad (55)$$

where

$$\hat{\xi}_- = \min(0, \xi_-), \quad \hat{\xi}_+ = \max(0, \xi_+) . \quad (56)$$

We use the flux terms (55) to evolve the hydrodynamic quantities within our Method of Lines scheme.

Our calculations of the wave speeds is approximate, following the methods outlined in [44] to increase the speed of the calculation at the cost of increased diffusivity. The method overstates the true wavespeeds by no more than a factor of $\sqrt{2}$, and then only for certain magnetic field and fluid velocity configurations. When computing the wavespeeds, we replace the full MHD dispersion relation by the approximate quadratic form (see 27 and 28 of [44]),

$$\omega_d^2 = k_d^2 \left[v_A^2 + c_s^2 \left(1 - \frac{v_A^2}{c^2} \right) \right] , \quad (57)$$

where we define the wave vector $k_\mu \equiv (-\omega, k_i)$, the fluid sound speed c_s , the Alfvén velocity,

$$v_A \equiv \sqrt{\frac{b^2}{\rho h + b^2}} = \sqrt{\frac{b^2}{\rho h^*}} , \quad (58)$$

the projected wave vector

$$K_\mu \equiv (\delta_\mu^\nu + u_\mu u^\nu) k_\nu , \quad (59)$$

and the dispersion relation between frequency and (squared) wave number as

$$\omega_d = k_\mu u^\mu = -\omega u^0 + k u^i , \quad (60)$$

$$k_d^2 = K_\mu K^\mu = \omega_d^2 + g^{\nu\sigma} k_\nu k_\sigma . \quad (61)$$

The resulting quadratic may be written

$$\begin{aligned} \xi^2 [W^2(V^2 - 1) - V^2] - 2\xi [\alpha W^2 \tilde{v}^i (V^2 - 1) + V^2 \beta^i] + \\ [(\alpha W \tilde{v}^i)^2 (V^2 - 1) + V^2 (\alpha^2 \gamma^{ii} - \beta^i \beta^i)] = 0 , \end{aligned} \quad (62)$$

where $V^2 \equiv v_A^2 + c_s^2(1 - v_A^2)$ and ξ is the resulting wavespeed. Note that the indices are not to be summed over. Instead, we find different wavespeeds in different directions.

When divergence cleaning is used to dissipate spurious numerical constraint violations that appear in the magnetic field, the characteristic wavespeeds include two additional (*luminal*) modes of the divergence cleaning subsystem. Since the divergence cleaning subsystem decouples from the remainder of the evolution system, its wavespeeds must be handled slightly differently [88, 89], which we discuss below in section 4.5.1.

4.4. Conservative to Primitive variable transformations

In GRHD, converting the conservative variables back to the primitives is a relatively straightforward task, which can be accomplished by inverting (10) – (12) with all of the magnetic field quantities set to zero. GRHydro accomplishes the task through a 1D Newton-Raphson scheme (summarised in section 5.5.4 of [22] and described in detail in the code documentation [90]). The scheme iterates by estimating the fluid pressure, determining the density, internal energy and the conservative quantities given this estimate, and using those in turn to calculate a new value for the pressure and its residual. The method works for any EOS, so long as one can calculate the thermodynamic derivatives $dP/d\rho|_\epsilon$ and $dP/d\epsilon|_\rho$.

MHD adds several complications to the inversion, although the additional equation, (13), is immediately invertible, yielding $B^i = \mathcal{B}^i/\sqrt{\gamma}$. Instead, the primary difficulty is that b^μ cannot be immediately determined, and there is no simple analogue to the GRHD expressions that allow us to calculate the density easily once the pressure is specified. As a reminder, if we consider the (known) values of the undensitised conservative variables,

$$\begin{aligned} \hat{D} &\equiv \frac{D}{\sqrt{\gamma}} = \rho W , & \hat{S}^i &\equiv \frac{S^i}{\sqrt{\gamma}} = \rho h^* W^2 v^i - \alpha b^0 b^i , \\ \hat{\tau} &\equiv \frac{\tau}{\sqrt{\gamma}} = \rho h^* W^2 - P^* - (\alpha b^0)^2 - \hat{D} , \end{aligned} \quad (63)$$

the GRHD system (in which $P^* = P$, $h^* = h$, and $b^\mu = 0$) allows us to define

$$Q \equiv \hat{\tau} + \hat{D} + P = \rho h W^2 , \quad (64)$$

and then determine the density as a function of pressure through the relation

$$\rho = \frac{\hat{D} \sqrt{Q^2 - \gamma_{ij} \hat{S}^i \hat{S}^j}}{Q} , \quad (65)$$

and thus the Lorentz factor and internal energy as well (see, e.g., [90]). In GRMHD, the most efficient approach for inverting the conservative variable set is often to use a multi-dimensional Newton-Raphson solver, with simplifications possible for barotropic EOS, for which the internal energy is assumed to be a function of the density only, eliminating the need to evolve the energy equation. Our methods follow very closely those of [63], particularly the 2D and 1D_W solvers they discuss.

We define a few auxiliary quantities for use in our numerical calculations. Unfortunately, it is impossible to construct a consistent notation that agrees with both the Einstein Toolkit release paper [22] and the paper that describes the conservative to primitive variable inversion scheme [63], so we choose to be consistent with the former here, noting the key differences in Appendix A.

From the values of the conservative variables and the metric, we construct the momentum density

$$\mathcal{S}_\mu \equiv -n_\nu T^\nu{}_\mu = \alpha T_\mu^0 , \quad (66)$$

whose spatial components are given by the relation $\mathcal{S}_i \equiv \hat{S}_i$ and its normal projection $\tilde{\mathcal{S}}$ given by

$$\tilde{\mathcal{S}}_\mu = (\delta_\mu^\nu + n_\mu n^\nu) \mathcal{S}_\nu . \quad (67)$$

Inside the Newton-Raphson scheme we make use of an auxiliary variable Q defined analogously to (64) by the expression

$$Q \equiv \rho h W^2 . \quad (68)$$

Since the inversion methods for general EOS and barotropic ones differ in some of the details, we discuss each in turn.

4.4.1. General EOS The 2D Newton-Raphson approach implemented solves the following two equations for the unknown quantities Q and v^2 , with all other terms known from the given conserved set:

$$\tilde{\mathcal{S}}^2 = \hat{S}_i \hat{S}^i = v^2 (B^2 + Q)^2 - (\mathcal{S} \cdot \mathbf{B})^2 \frac{B^2 + 2Q}{Q^2} , \quad (69)$$

$$\mathcal{S} \cdot n = -(\hat{\tau} + \hat{D}) = -\frac{B^2}{2}(1 + v^2) + \frac{(\mathcal{S} \cdot \mathbf{B})^2}{2} Q^{-2} - Q + P , \quad (70)$$

where all dot products are understood as four-dimensional when involving four-vectors and three-dimensional when not: e.g., $\mathcal{S} \cdot \mathbf{B} = \mathcal{S}_\mu B^\mu$, which is equivalent to $\mathcal{S}_i B^i$ since B^μ is a three-vector and thus $B^0 = 0$.

For a polytropic/Gamma-law EOS where $P = (\Gamma - 1)u$ and $u \equiv \rho\epsilon$ is the internal energy density, we may calculate P from the conserved variables and the current guess for Q and v^2 by noting that

$$P = \frac{\Gamma - 1}{\Gamma} \left[(1 - v^2)Q - \hat{D}\sqrt{1 - v^2} \right] . \quad (71)$$

The iteration updates Q and v^2 subject to the consistency conditions that $0 \leq v^2 \leq 1$ and $Q > 0$, and a post-iteration check is performed to ensure that $\epsilon > 0$.

For a more general EOS where (71) does not apply, and given the structure of the EOS interface in the Einstein Toolkit, it is easier to solve for the internal energy density u , and use it as a variable in a three-dimensional Newton scheme along with Q and v^2 . First, noting that

$$P = (1 - v^2)Q - \hat{D}\sqrt{1 - v^2} - u , \quad (72)$$

we may rewrite (70) as

$$\mathcal{S} \cdot \mathbf{n} = -\frac{B^2}{2}(1 + v^2) + \frac{(\mathcal{S} \cdot B)^2}{2}Q^{-2} - v^2Q - \hat{D}\sqrt{1 - v^2} - u . \quad (73)$$

In the Newton-Raphson steps, we first set $\rho = \hat{D}\sqrt{1 - v^2}$ and solve

$$u + P(\rho, u) = Q(1 - v^2) - \hat{D}\sqrt{1 - v^2} , \quad (74)$$

where the left-hand side will, in general, be a monotonic function of ρ and u , along with (69) and (73).

Noting that the density depends only on the constant \hat{D} and the value of v^2 within the Newton-Raphson scheme, the partial derivatives of the pressure with respect to the Newton-Raphson variables are given by

$$\frac{\partial P}{\partial(v^2)} = \left(\frac{\partial P}{\partial \rho} \right)_u \frac{d\rho}{d(v^2)} = \left[\left(\frac{\partial P}{\partial \rho} \right)_\epsilon - \frac{\epsilon}{\rho} \left(\frac{\partial P}{\partial \epsilon} \right)_\rho \right] \frac{d\rho}{d(v^2)} , \quad (75)$$

$$\frac{\partial P}{\partial u} = \frac{1}{\rho} \left(\frac{\partial P}{\partial \epsilon} \right)_\rho . \quad (76)$$

We write these terms in this way, since the Einstein Toolkit EOS interface uses the pressure as a function of density and the *specific* internal energy ϵ , $P = P(\rho, \epsilon)$, and calculates partial derivatives against those variables, rather than ρ and u .

4.4.2. Barotropic EOS For cases where the pressure and internal energy are functions of the rest mass density only, the EOS is barotropic (a polytropic $P = K\rho^\Gamma$ EOS is a special case of a barotropic EOS), we need a different inversion technique, since the quantity $\mathcal{S} \cdot \mathbf{n}$ in (70) requires knowledge of $\hat{\tau}$, which is not evolved in these cases. Instead, the inversion uses (69) only, using it to eliminate the variable v^2 from the Newton-Raphson scheme by solving for $v^2(Q)$:

$$v^2(Q) = \frac{Q^2\tilde{\mathcal{S}}^2 + (\mathcal{S} \cdot B)^2(B^2 + 2Q)}{Q^2(B^2 + Q)^2} . \quad (77)$$

In the special case of a polytropic EOS, we proceed by first solving for $\rho(Q)$ through an independent Newton-Raphson loop over the equation

$$\rho Q = \hat{D}^2 \left(1 + \frac{\Gamma K \rho^{\Gamma-1}}{\Gamma - 1} \right) . \quad (78)$$

Next, we use (77) and the fact that

$$v^2 = \frac{\rho^2}{\hat{D}^2} - 1 , \quad (79)$$

to replace (69) by an expression given only in terms of Q and $\rho(Q)$:

$$\begin{aligned} 0 &= Q^2(B^2 + Q)^2 v^2 - Q^2(B^2 + Q)^2 v^2 , \\ &= Q^2 \tilde{\mathcal{S}}^2 + (\mathcal{S} \cdot B)^2 (B^2 + 2Q) - \left(\frac{\rho^2}{\hat{D}^2} - 1 \right) Q^2 (B^2 + Q)^2 . \end{aligned} \quad (80)$$

When performing the Newton-Raphson step, all quantities in (80) are known a priori except the iteration variable Q and $\rho(Q)$, which depends upon it. The loop over Q makes use of the derivative of (78), given by

$$\frac{d\rho}{dQ} = \frac{\rho}{\hat{D}^2 \gamma K \rho^{\Gamma-2} - Q} . \quad (81)$$

4.5. Divergence-free Constraint Treatment

One of the main difficulties in numerical MHD simulations is the appearance of non-zero divergence of the magnetic field due to numerical errors, violations of (18) that would be interpreted as magnetic monopoles. A number of techniques have been designed to combat these, including “divergence cleaning” [41, 42, 91], a method which introduces a new field that both damps and advects divergences off the grid, “constrained transport” algorithms [40] that balance out fluxes exactly to maintain divergence-free magnetic fields to round-off accuracy, and magnetic vector potential methods that seek to achieve the same result [69, 92]. Constrained transport methods are often difficult to implement in simulations that employ mesh refinement, since maintaining balance at refinement boundaries is algorithmically complex. In the initial GRHydro MHD release we implement both divergence cleaning and constrained transport, both of which are detailed below.

4.5.1. Divergence cleaning Divergence cleaning works by introducing a new field variable that both damps divergences and advects them off the grid through a hyperbolic equation modelled after the telegraph equation, driving numerical solutions towards zero divergence. Our implementation follows closely that of [41, 42], with some minor differences. The new field variable ψ satisfies the evolution equation

$$\nabla_\mu \left({}^*F^{\mu\nu} + g^{\mu\nu} \psi \right) = \kappa n^\nu \psi , \quad (82)$$

a modification of Maxwell’s equations (9) that reduces to the familiar form as $\psi \rightarrow 0$. Our parameter κ , as we note below, determines the damping rate of the divergence

cleaning field ψ , and incorporates the ratio of the parabolic to the hyperbolic damping speed dependence for the scheme, often denoted respectively by c_p and c_h in other works.

The evolution equation for ψ is given by the time component of (82). Noting that [73]

$${}^*F^{\mu\nu} = \frac{1}{W} (u^\mu B^\nu - u^\nu B^\mu) , \quad (83)$$

the first term on the left-hand-side of (82) yields, after some algebra and use of the fact that $B^0 = -\alpha F^{00} = 0$,

$$\nabla_\mu {}^*F^{\mu 0} = -\frac{1}{\alpha\sqrt{\gamma}} \partial_i \sqrt{\gamma} B^i , \quad (84)$$

and the second term

$$\nabla_\mu g^{\mu 0} \psi = g^{\mu 0} \partial_\mu \psi = \frac{1}{\alpha^2} \left[-\partial_t \psi + \beta^i \partial_i \psi \right] . \quad (85)$$

Combining the two, we find

$$\partial_t \psi + \partial_i \left(\alpha B^i - \psi \beta^i \right) = \psi \left(-\kappa \alpha - \partial_i \beta^i \right) + \sqrt{\gamma} B^i \partial_i \left(\frac{\alpha}{\sqrt{\gamma}} \right) , \quad (86)$$

where we have grouped the derivative terms to respect the flux-conservative form.

The spatial part of (82) is the modified evolution equation for the magnetic field. The two terms on the left hand side yield, respectively,

$$\nabla_\mu {}^*F^{\mu j} = \frac{1}{\alpha\sqrt{\gamma}} \left\{ \partial_t \sqrt{\gamma} B^j + \partial_i \sqrt{\gamma} \left[\left(\alpha v^i - \beta^i \right) B^j - \left(\alpha v^j - \beta^j \right) B^i \right] \right\} , \quad (87)$$

$$\nabla_\mu g^{\mu j} \psi = g^{\mu j} \partial_\mu \psi = \frac{1}{\alpha^2} \left[\beta^j \partial_t \psi + \left(\alpha^2 \gamma^{ij} - \beta^i \beta^j \right) \partial_i \psi \right] . \quad (88)$$

Using (86) to eliminate the $\partial_t \psi$ term, switching over to \mathcal{B}^i as the primary magnetic field variable, and combining derivatives to produce a flux-conservative form, we find after some algebra,

$$\begin{aligned} \partial_t \mathcal{B}^j + \partial_i \left[\left(\alpha v^i - \beta^i \right) \mathcal{B}^j - \alpha v^j \mathcal{B}^i + \alpha \sqrt{\gamma} \gamma^{ij} \psi \right] = \\ - \mathcal{B}^i \partial_i \beta^j + \psi \partial_i \left(\alpha \sqrt{\gamma} \gamma^{ij} \right) , \end{aligned} \quad (89)$$

which reduces to the standard evolution equation for $\psi \rightarrow 0$ and $\partial_i \mathcal{B}^i \rightarrow 0$. To evaluate the right-hand-side, we make use of the identity

$$\partial_i (\sqrt{\gamma} \gamma^{ij}) = -\sqrt{\gamma} \gamma^{kl} \Gamma_{kl}^j = \sqrt{\gamma} \left[\frac{1}{2} \gamma^{ij} \gamma^{kl} \partial_i \gamma_{kl} - \gamma^{jk} \gamma^{il} \partial_i \gamma_{kl} \right] . \quad (90)$$

The characteristic velocities of the evolution system for MHD without and with divergence cleaning differ. In the former, the largest-magnitude wave speed is the fast magnetosonic wave speed. The inclusion of divergence cleaning introduces two additional modes, corresponding to eigenvalues for the evolution of the divergence cleaning field ψ and the longitudinal component of the magnetic field, (i.e., the case where $i = j$ inside the term in brackets in (90)), that in the flat spacetime case decouple from the remaining seven eigenvalues of the system [88, 89]. In our scheme, both modes have characteristic speed equal to the speed of light. Our system corresponds to setting

the hyperbolic damping speed c_h of the scheme to be the speed of light in the notation of Newtonian divergence cleaning methods [91]. This is in accordance with standard practice in both Newtonian and relativistic calculations to choose a characteristic speed for divergence cleaning that is as fast as allowed without violating causality, or in the case of Newtonian calculations, the Courant condition. We note that for the luminal-speed modes, the HLLE flux formula, (55) reduces to the local Lax-Friedrichs form when evaluated on a Minkowski background,

$$F(U) = \frac{1}{2} \left[F(U^L) + F(U^R) - \hat{\xi}(U^R - U^L) \right], \quad (91)$$

where $\hat{\xi} = \max(\hat{\xi}_+, |\hat{\xi}_-|)$. Test runs performed with the HLLE Riemann solver, but without the luminal-speed velocities often develop large spurious oscillations, particularly in cases where strong rarefactions are present, e.g., the cylindrical blast wave evolution described in section 5.3 and the rotor test described in section 5.4.

In our current implementation, we use the hyperbolic divergence cleaning technique only for flat spacetime tests. In this case the divergence cleaning subsystem decouples from the remainder of the MHD evolution system. In curved spacetime this decoupling is not as obvious and we opt to use constrained transport techniques instead, which we describe in the following.

4.5.2. Constrained transport As an alternative to divergence cleaning, and to simplify comparison of the performance of the new code with existing GRMHD codes, we also implemented a variant of a constrained transport scheme [93]. In constrained transport schemes one carefully constructs a numerical update scheme such that the divergence free constraint (19) is conserved to numerical round-off accuracy. Rather than implementing the original scheme proposed in [93–95], which relies on a complicated staggering of the magnetic field components, we employ the simplified scheme called “flux-CT” described in [40, 43, 83], a two-dimensional version of which is used in the **HARM** code [44]. The scheme uses cell-centred values of the magnetic field. For completeness of presentation we reproduce the basic description of the scheme found in [43], but refer the reader to the original literature for more details.

In constrained transport schemes the induction equation (90) is written in terms of the electric field $\vec{\mathcal{E}}$ at the edges of each face of a simulation cell (see figure 1).

Employing Ampère’s law, the time derivative of the face-averaged magnetic field component $\hat{\mathcal{B}}^x$ is given by

$$\begin{aligned} \frac{\partial \hat{\mathcal{B}}_{i+\frac{1}{2},j,k}^x}{\partial t} \Delta y \Delta z = & -\mathcal{E}_{i+\frac{1}{2},j,k-\frac{1}{2}}^y \Delta y - \mathcal{E}_{i+\frac{1}{2},j+\frac{1}{2},k}^z \Delta z + \mathcal{E}_{i+\frac{1}{2},j,k-\frac{1}{2}}^y \Delta y \\ & + \mathcal{E}_{i+\frac{1}{2},j-\frac{1}{2},k}^z \Delta z, \end{aligned} \quad (92)$$

with analogous equations for the other magnetic field components. In the ideal MHD approximation, the electric field components in (93) can be expressed in terms of the fluxes $\mathcal{B}^k \tilde{v}^i - \mathcal{B}^i \tilde{v}^k$ for the magnetic field $\vec{\mathcal{B}}$ given in (16). Specifically, we use the numerical fluxes of the induction equation to calculate these electric field components.




Figure 1. Sketch of a simulation cell showing the location of face centred magnetic field and edge centred electric fields.

Since the scheme of [43] evolves the cell-averaged magnetic field values, it computes the change in the cell-centred fields as the average change in the face centred fields

$$\frac{\partial \mathcal{B}_{i,j,k}^x}{\partial t} = \frac{1}{2} \left(\frac{\partial \hat{\mathcal{B}}_{i-\frac{1}{2},j,k}^x}{\partial t} + \frac{\partial \hat{\mathcal{B}}_{i+\frac{1}{2},j,k}^x}{\partial t} \right), \quad (93)$$

which is the actual evolution equation implemented in the code.

In this formalism, the conserved divergence operator is given by

$$\begin{aligned} (\nabla \cdot \mathcal{B})_{i+\frac{1}{2},j+\frac{1}{2},k+\frac{1}{2}} &= \frac{1}{4\Delta x} \sum_{j'=j}^{j+1} \sum_{k'=k}^{k+1} (\mathcal{B}_{i+1,j',k'}^x - \mathcal{B}_{i,j',k'}^x) \\ &+ \frac{1}{4\Delta y} \sum_{i'=i}^{i+1} \sum_{k'=k}^{k+1} (\mathcal{B}_{i',j+1,k'}^y - \mathcal{B}_{i',j,k'}^y) \\ &+ \frac{1}{4\Delta z} \sum_{i'=i}^{i+1} \sum_{j'=j}^{j+1} (\mathcal{B}_{i',j',k+1}^z - \mathcal{B}_{i',j',k}^z), \end{aligned} \quad (94)$$

which can be interpreted as a cell corner-centred definition of a divergence operator.

5. Tests

5.1. Monopole tests

In order to test the divergence cleaning formalism and to explore its properties, we implement a number of tests that initialise the MHD configuration with numerical monopoles. Each case uses a uniform-density, uniform-pressure fluid which is initially at rest, and we assume Minkowski spacetime. The ambient magnetic field is assumed to be zero. The initial magnetic configurations include “point” monopoles, where $B^x \neq 0$ is set at a single point in the Center of the domain, and Gaussian monopoles for which

$$B^x = \begin{cases} e^{-r^2/R_G^2} - e^{-1}; & r < R_G, \\ 0; & r \geq R_G, \end{cases} \quad (95)$$

where R_G is the radius of the compactly supported monopole. In the supplementary material we present additional results on the performance of the algorithm when dealing with high frequency constraint violations.

In Figs. 2, we show the evolution of Gaussian and 3-dimensional alternating Gaussian monopole initial data. In each case, the grid is chosen to be 200^3 , spanning a coordinate range from -2 to 2 in each dimension, and we set $R_G = 0.2$. The simulations are performed using second-order Runge-Kutta (RK2) time integration and TVD-based reconstruction with an monotonized central limiter and a Courant-Friedrichs-Lewy (CFL) factor of 0.25 (i.e., $\Delta t/\Delta x = 0.25$ here). The fluid, assumed to follow a $\Gamma = 5/3$ ideal-gas law, i.e. to follow the condition $P = (\Gamma - 1)\rho\epsilon$, was set to an initially stationary state with density and internal energy given by $\rho = 1.0$ and $\epsilon = 0.1$, respectively. We vary the divergence cleaning parameter κ that appears in the driving term in (82), choosing values $\kappa = 1$, 10 , and 100 .

There is a markedly different evolution of the divergence of the magnetic field in time, following a predictable pattern.

The choice $\kappa = 1$ yields the slowest damping rate, which allows the wave-like behaviour of the divergence cleaning field to radiate away the divergence of the B -field away from its original location as it damps downward. By contrast, $\kappa = 100$ yields a much stiffer system of equations, resulting in a relatively slow damping rate, particularly for the lower-frequency terms in the initial data, with very low divergence in the magnetic field spreading through the numerical grid. The intermediate case, $\kappa = 10$, yields the closest analogue to a critically damped system in that the amplitude of the divergence decreases most rapidly, while errors are efficiently radiated away across the grid. In general, we would recommend values of $\kappa \sim 10$ to be used as a default.

5.2. Planar MHD Shocktubes

Historically, the simplest tests for an MHD scheme are shocktubes on a static flat-space Minkowski metric. They are important to demonstrate the ability of the new code to capture a variety of MHD wave structures. Despite being simple setups they do provide a stringent tests for the algorithm. For our one-dimensional shocktube tests, we set up “Left” and “Right” MHD states on either side of a planar interface, with parameters drawn from the five cases considered in [96] and summarised in Table 1. These have been widely used by several groups to establish the validity of numerical MHD codes. These parameter choices include generalisations of familiar cases long used to test non-relativistic MHD codes, e.g., case “Balsara1”, which is a generalisation of the Brio-Wu shock tube problem [97].

To validate the new code, we evolve each of these shocks in each of the coordinate directions (i.e., “Left” states for domains of the grid satisfying $x < 0$, $y < 0$, and $z < 0$ and “Right” states for $x > 0$, $y > 0$, and $z > 0$, respectively), finding excellent agreement (limited only by numerical precision) in each direction, as expected. We also evolve shocks with mid-planes oriented obliquely to the coordinate planes, choosing




Figure 2. Behaviour of our divergence cleaning scheme, demonstrated by monopole damping and advection for a magnetic field with an initially Gaussian B^x profile of radius $R_G = 0.2$. We show results for different values of the divergence cleaning damping parameter κ . For $\kappa = 100$ (dot-dashed blue curve), the system is quite stiff and damping is substantially slower than in other cases and more of the divergence remains localised near the initial position, while for $\kappa = 1$ (dashed red curve) the damping is only slightly faster but spreads more rapidly away across the grid. For $\kappa = 10$ (solid green curve), we find a nearly ideal choice of rapid damping and advection of the error away from the source, and recommend this value as a default for general calculations.

$x + y = 0$ (“2D diagonal”) and $x + y + z = 0$ (“3-d diagonal”) as the mid-planes. Whereas shocks in any of the coordinate planes initially satisfy the divergence free constraint to roundoff error, and maintain this state indefinitely given the symmetry in the setup, diagonal shocks yield non-zero divergence of the magnetic field across the shock front due to finite resolution effects. The simplest way to understand this is to realise that for a 1-dimensional domain (say along the x direction), the x component of the magnetic field B^x is necessarily constant in space and time as a consequence of

the induction equation (90). The slab symmetry of the test on the other hand implies that all fields, in particular the B^y and B^z field, are independent of y and z . These two effects imply that $\partial_i B^i = 0$ to roundoff precision initially and for all times, since the finite volume scheme preserves this property. Diagonal shocks, on the other hand, break this symmetry by having field components depend on two coordinates introducing divergence constraint violations of order of the truncation error. For non-flat geometries, the non-linear coupling of GRMHD equations would generate violations on the order of the truncation error in all directions. This multidimensional test allows us to verify that these constraint violations are controlled by employing divergence cleaning.

In all cases, the coordinate origin was chosen so that no grid cell was centred directly on the shock front, so that each cell was unambiguously located in the left or right states. In each run we used RK2 time integration with a CFL factor of 0.8 for 1-d shocks and 0.4 for 2D diagonal cases, with $\Delta x = 1/1600$ and $\Delta x = \Delta y = 1/(800\sqrt{2})$ for the 1-d and 2D diagonal cases to ensure equal resolution in the shock direction. For all results shown below, we use the HLLE Riemann solver and TVD reconstruction of the primitive variables with a monotonized central limiter. Boundary conditions are implemented by “copying” all data to grid edges from the nearest points in the interior located on planes parallel to the shock front. Divergence cleaning is turned on for the 2D case with $\kappa = 10$. Table 1 lists the parameters describing each shock test that is performed, and we note that all of them assume an ideal-gas law EOS with the given value of Γ , here $\Gamma = 2$ for the “Balsara1” shock test and $\Gamma = 5/3$ for the others. These tests are also included as options in the code release within the `GRHydro_InitData` thorn. Exact solutions for each case are computed using the open-source available code of [98, 99]. In all cases, we see very good agreement with the expected results that are comparable to results of similar codes in the literature [43, 66, 70, 73, 96, 98, 100, 101]. We conclude that divergence cleaning does not interfere with generating physically meaningful GRMHD evolution results in the presence of shocks, and that the new code performs as well as other codes based on the same numerical techniques [89].

In figure 3 we show results for the relativistic generalisation of the Brio & Wu shock test develop in [96]. The initial shock develops into a left-going fast rarefaction, a left-going compound wave, a contact discontinuity, a right-going slow shock and a right-going fast rarefaction. We find that the code captures all elementary waves and is in good agreement with the exact solution of [43]. In panel (g) of 3 we show the constraint violation as measure by a 2nd order centred finite difference stencil for $\partial_i \mathcal{B}^i$. Figure 4 displays the results of the relativistic MHD collision problem (problem 4 of [96] in which two very relativistic (Lorentz factor 22.37) streams collide with each other. We find deviations from the exact solution that are on the same level as in [43, 96].

5.3. Cylindrical Shocks

A more stringent multidimensional code test is provided by a cylindrical blast wave expanding outward in two dimensions. We take the parameters for this test problem




Figure 3. Evolution of the Balsara1 shocktube, performed with divergence cleaning, showing 1D and 2D diagonal cases as solid (blue) and dashed (red) curves, respectively. From top to bottom, we show the density ρ , gas pressure P , normal and tangential velocities v^x and v^y , the Lorentz factor W , the tangential magnetic field component B^y , and the numerical divergence of the magnetic field for the 2D case (for 1D shocks, the divergence is uniformly zero by construction). Parameters for the initial shock configuration are given in Table 1. All results are presented at $t = T_{\text{ref}}$, given by the final column of Table 1. The results agree well with those reported in [43, 96] with most plots being indistinguishable and only the Lorentz factor plot showing a slight overshoot on the left-hand side of the shock. A more detailed description of the test setup and parameters can be found in the main text in section 5.2.




Figure 4. Evolution of the Balsara4 shocktube case, performed with divergence cleaning, with all conventions as in figure 3. The results code reproduces the exact result very well, with deviations from the exact solution similar to those found in [96]. A more detailed description of the test setup and parameters can be found in the main text in section 5.2.

Table 1. Shock tube parameters. For each of these shock tube tests, originally compiled in [96], we list for both the left “ L ” and right “ R ” states the fluid density ρ , specific internal energy ϵ , fluid 3-velocity $\vec{v} \equiv v^i$, and tangential magnetic field B^t , as well as the (uniform) normal magnetic field magnitude B^n , the adiabatic index Γ of the ideal-gas law equation of state, and the time T_{ref} at which results are plotted. Our values are quoted for shocks moving in the x -direction, i.e., the shock front is oriented in the $y - z$ plane, so B^t is equivalent to (B^y, B^z) and $B^n = B^x$. For diagonal shocks, all quantities are rotated appropriately along the diagonal of $x - y$ plane. Results for cases marked by an asterisk* are shown in the supplementary material only.

Name	ρ_L	ϵ_L	\vec{v}_L	B_L^t	ρ_R	ϵ_R	\vec{v}_R
Balsara1	1.0	1.0	$\vec{0}$	(1.0,0)	0.125	0.8	$\vec{0}$
Balsara2*	1.0	45.0	$\vec{0}$	(6.0,6.0)	1.0	1.5	$\vec{0}$
Balsara3*	1.0	1500.0	$\vec{0}$	(7.0,7.0)	1.0	0.15	$\vec{0}$
Balsara4	1.0	0.15	(0.999, 0, 0)	(7.0,7.0)	1.0	0.15	(-0.999,0,0)
Balsara5*	1.08	1.425	(0.4, 0.3, 0.2)	(0.3,0.3)	1.0	1.5	(-0.45,-0.2,0.2)

Name	B_R^t	B^n	Γ	T_{ref}
Balsara1 (continued)	(-1.0,0)	0.5	2	0.4
Balsara2* (continued)	(0.7,0.7)	5.0	5/3	0.4
Balsara3* (continued)	(0.7,0.7)	10.0	5/3	0.4
Balsara4 (continued)	(-7.0,-7.0)	10.0	5/3	0.4
Balsara5* (continued)	(-0.7,0.5)	2.0	5/3	0.55

from [102]. The density profile is determined by two radial parameters, r_{in} and r_{out} , and two reference states, such that

$$\rho(r) = \begin{cases} \rho_{\text{in}} & ; \quad r \leq r_{\text{in}} , \\ \exp \left[\frac{(r_{\text{out}} - R) \ln \rho_{\text{out}} + (r - r_{\text{in}}) \ln \rho_{\text{in}}}{r_{\text{out}} - r_{\text{in}}} \right] & ; \quad r_{\text{in}} < r < r_{\text{out}} , \\ \rho_{\text{out}} & ; \quad r \geq r_{\text{out}} , \end{cases} \quad (96)$$

with an equivalent form for the pressure gradient. The initial fluid velocity is set to zero, and the initial magnetic is uniform in the domain. We use the following shock parameters:

$$\begin{aligned} r_{\text{in}} &= 0.8, \quad r_{\text{out}} = 1.0; \quad \rho_{\text{in}} = 10^{-2}, \quad \rho_{\text{out}} = 10^{-4}; \quad P_{\text{in}} = 1.0, \\ P_{\text{out}} &= 3 \times 10^{-5}; \quad B^i = (0.1, 0, 0) . \end{aligned} \quad (97)$$

All tests use Γ -law equation of state with adiabatic index $\Gamma = 4/3$. We use a $200 \times 200 \times 8$ grid, spanning the coordinate range $[-6, 6]$ in the x and y -directions, which is the setup used in [102]. To verify that our results are indeed convergent with resolution and to compare to [69], we run two more simulations using 400 and 500 points as well. All tests shown use the divergence cleaning scheme with $\kappa = 5$, HLLE Riemann solver, TVD reconstruction using a monotonized central limiter. No explicit dissipation is added to the system of equations.

This test is known to push the limits of many combinations of MHD evolution techniques, and often fails for specific combinations of reconstruction methods and

Riemann solvers [67, 69]. In our tests we only explore a small range of settings: we verify that the test succeeds when we swap TVD reconstruction with a 2nd order ENO scheme or use constraint transport instead of divergence cleaning. In the later case we add Kreiss-Oliger dissipation [103] of order 3 and strength parameter $\epsilon = 3$ [103] to the magnetic field variables to stabilise the system.

We find this test to be among the most sensitive one of our tests. Including the two light-like modes of the divergence cleaning field (see section 4.5.1) is crucial to obtain correct results. When these are not included, the code crashes quickly due to the growth of spurious oscillations in the magnetic field in the rarefaction region that develops behind the shock, particularly along the diagonals where the shock front expands obliquely to the grid (as was seen in tests of a several other codes [66, 67, 69]). We note that implementing the local Lax-Friedrichs flux (see (91)) instead of HLLE also stabilises the evolution.

Two-dimensional profiles are shown in figure 5 for P , W_{Lorentz} , B^x , and B^y , along with numerically determined magnetic field lines, at $t = 4$. These results agree well with those shown in figure 10 of [102]. One-dimensional slices along the x - and y -axes for the rest mass density, gas pressure, magnetic pressure and Lorentz factor at $t = 4$ are shown in figure 6 for three different numerical resolutions. Our results show good qualitative agreement with those presented previously by [67] and particularly by [69], once one accounts for rescalings associated with different choices for the initial shock parameters. We find that the benefit of increasing the resolution is largest in the region within the shock front, which is consistent with results obtained by other groups [67, 69].

5.4. Magnetic rotor

As a second two-dimensional test problem we simulate the magnetic rotor test first described in [95] for classical MHD and later generalised to relativistic MHD in [101]. The setup consists of a cylindrical column (the *rotor*) of radius $r_{\text{in}} = 0.1$ and density $\rho_{\text{in}} = 10$ embedded in a medium of lower density $\rho_{\text{out}} = 1.0$. Initially the pressure inside the column and in the medium are equal $P_{\text{in}} = P_{\text{out}} = 1$ and the cylinder rotates with uniform angular velocity of $\Omega = 9.95$ along the cylinder axis, so that the fluid 3-velocity reaches a maximum value of $v_{\text{max}} = 0.995$ at the outer edge of the cylinder. The fluid outside the cylinder is initially at rest. The equation of state used is a Γ -law, with $\Gamma = 5/3$. The cylinder is threaded by an initially uniform magnetic field of magnitude $B^x = 1.0$ along the x direction, covering the entire space; cylinder and exterior. We use TVD reconstruction with the minmod limiter, the HLLE Riemann solver, RK2 time stepping with CFL factor 0.25. The magnetic field is evolved using the divergence cleaning technique using a damping factor $\kappa = 5.0$.

At the beginning of the simulation, a strong discontinuity is present at the edge of the cylinder since we do not apply any smoothing there. During the simulation, magnetic braking slows down the rotor while the magnetic field lines themselves are dragged from their initial horizontal orientation. At the end of the simulation, at $t = 0.4$, the field




Figure 5. Evolved state of the cylindrical explosion test problem, with parameters based on the test presented in [102]. At $t = 4$, we show the following quantities: top panels: Gas pressure P and Lorentz factor W , along with numerically determined magnetic field lines; bottom panels, B^x and B^y . The numerical resolution is $\Delta x = 0.06$. Profiles are almost indistinguishable from the solutions found in figure 10 of [102]. The very low pressure region outside of the shock front is slightly more extended in [102]. It is worth noting that [102] apply some numerical resistivity to control artifacts, while our divergence cleaning scheme has dissipative terms present in the divergence cleaning equations but do not add explicit resistivity, hence our results in low pressure regions might well differ slightly.




Figure 6. One-dimensional slices along the x - and y -axes for the evolved state of the cylindrical explosion test. Displayed are the rest mass density, gas pressure, magnetic pressure and Lorentz factor at $t = 4$. The red solid lines corresponds to a resolution of $\Delta x = 0.06$, the black dashed one corresponds to a resolution of $\Delta x = 0.03$ and the blue dotted on corresponds to a resolution of $\Delta x = 0.024$. These results may be compared to figure 7 of [67] and particularly figure 7 of [69], noting that their parameters differ from ours, which were chosen to match the test as presented by [102], particularly by the presence of a narrow transition region whose effects are visible as a secondary shock in the outer edges of the shock front. Comparing the plots, both codes agree on the overall structure of the result, with the two-dimensional pseudocolor plots in 5 being indistinguishable and the one-dimensional slices clearly showing the same structure of the shocks. A more detailed discussion of the test setup and results can be found in the text.

lines in the central region have rotated by nearly 90 degrees while at large radii the orientation of the field lines remains unchanged.

Our results, shown in figure 7 for a grid spacing $\Delta x = 1/400$, compare well with those presented in figure 5 of [101] and figure 8 of [69]. The density profile at the end of simulation is reproduced, showing slightly less noisy behaviour at the location of the expanding high density shell than the results reported in [101]. In addition, the magnetic field lines displayed in figure 5 of [101] are also reproduced very well. A slight over-density near the left and right corners of the (now low density) rotor is present in our simulations that is absent in [101].

Figure 8 depicts one-dimensional slices through the simulation domain at the end of the run. We show results for three different resolution, using 250, 400 and 500 points to cover the domain. We find excellent agreement with previous work at similar resolutions [69].

5.5. Alfvén Wave

The propagation of low-amplitude, circularly-polarised Alfvén waves has an exact solution that is useful for testing the performance of the MHD sector of the code [40] and provides a scenario for clean convergence tests, since the solution is smooth. The initial configuration is a uniform-density fluid with a velocity profile given by

$$v^x = 0; \quad v^y = -v_A A_0 \cos(kx); \quad v^z = -v_A A_0 \sin(kx) , \quad (98)$$

and corresponding magnetic field configuration

$$B^x = \text{const.}; \quad B^y = A_0 B^x \cos(kx); \quad B^z = A_0 B^x \sin(kx) , \quad (99)$$

where we choose $B^x = 1.0$, an amplitude parameter $A_0 = 1.0$, $\rho = 1.0$. A Γ -law EOS with $\Gamma = 5/3$ was used, as well as TVD as reconstruction method, a 2nd order Runge-Kutta (RK) [104, 105] time integration with the Courant factor set to 0.2, and the HLLE Riemann solver, as well as divergence cleaning.

The Alfvén speed is given by the expression

$$v_A^2 = \frac{2(B^x)^2}{\rho h + (B^x)^2(1 + A_0^2)} \left[1 + \sqrt{1 - \left(\frac{2A_0(B^x)^2}{\rho h + (B^x)^2(1 + A_0^2)} \right)^2} \right]^{-1} , \quad (100)$$

and the wavevector $k = 2\pi/L_x$ for one-dimensional cases. For two-dimensional cases, we rotate the coordinates so that wave fronts lie along diagonals of the grid. In all cases, periodic boundary conditions are assumed. Several values for the pressure have been used for testing purposes by other groups, spanning the range $P = 0.1$ [40] to $P = 1.0$ [106], but here we choose $P = 0.5$, which yields the convenient result $v_A = 0.5$. In time, we expect the wave to propagate across the grid, such that if it travels an integer number of wavelengths we should, in the ideal case, reproduce our initial data exactly. In figure 9, we show the convergence results for our code for both 1-dimensional and 2-dimensional cases, finding the expected second-order convergence with numerical resolution when plotting the L_2 norm of the difference between the evolved solution and




Figure 7. Evolved state of the magnetic rotor test problem, with parameters based on the test presented in [101]. At $t = 0.4$, we show the following quantities: top panels: Density ρ and pressure P ; bottom panels, magnetic pressure $P_m \equiv \frac{b^2}{2}$ and Lorentz factor W , along with numerically determined magnetic field lines. Our results are in close agreement with those shown in [67, 69, 101].

the initial data, evaluated along the x -axis (a full two-dimensional comparison produces the same result). We note that while the overall error magnitude appears larger than that shown for one-dimensional and two-dimensional waves in similar works, e.g., [106], our results are shown after five periods, and we find comparable accuracy to [106] for the first cycle.

5.6. Loop advection

In this test, originally proposed in [107] and presented in a slightly modified form by [64, 106, 108], a region with a circular cross section is given a non-zero azimuthal




Figure 8. Evolved state of the magnetic rotor test problem, after $t = 0.4$. We display we display data for three different resolutions using 250, 400, and 500 cells each to cover the domain. The data was extracted along the $x = 0.5$ and $y = 0.5$ axis respectively. Our results are in very good agreement with figure 9 of [69].

magnetic field of constant magnitude (the “field loop”). The magnetic field is initialised to zero outside this region. The entire fluid configuration is given a uniform velocity and periodic boundary conditions are imposed. The analytical solution is for the field loop to move with constant velocity across the grid, with no magnetic field evolution. With $r_c \equiv \sqrt{x^2 + y^2}$ as the cylindrical radius we have an initial magnetic field

$$B_x, B_y = \begin{cases} -A_{loop}y/r_c, & A_{loop}x/r_c; & r < R_{loop} , \\ 0; & & r \geq R_{loop} . \end{cases} \quad (101)$$

In practice, this configuration results from the choice of the vector potential $\vec{A} = (0, 0, \max[0, A_{loop}(R_{loop} - r_c)])$, though we note that we here set the B -field values explicitly in the initial data, rather than by finite differencing of the A -field, an option




Figure 9. Shown is the L_2 norm of the difference between the evolved solution and the initial data of the Alfvén wave test. Convergence is observed for the one-dimensional (1D, blue, dotted line) and two-dimensional (2D, red, dashed line) test, along with scaling expected for second-order convergence (solid line). In each case, as a function of N_x , the number of points in the x -direction, we plot the L_2 norm of the B -field along a line drawn in the x -direction through the centre of the computational domain, measuring the absolute difference between the initial and final values of B^i after 5 full wave cycles for both the one-dimensional and two-dimensional cases. The resolutions used here were (16, 32, 64, 128) and $(20 \times 15, 40 \times 30, 80 \times 60, 160 \times 120)$ points in 1D and 2D respectively.

also included in the code.

Here, following [106], we choose parameters $A_{loop} = 10^{-3}$, $\rho = 1.0$, $P = 3.0$ for the initial magnetic field and hydrodynamic configuration. We consider two different models for the direction of the field loop: a “two-dimensional” case in which it is oriented vertically, and a “three-dimensional” case for which we rotate the direction of the tube so that it is oriented along the $x - z$ face diagonal of our rectangular grid. For the initial velocity field of the two-dimensional configurations, we set $v^x = 1/2$, $v^y = 1/24$ and consider two different cases for the vertical velocity field parallel to the loop: one in which $v^z = 0$, so that the motion is perpendicular to the field loop, and one in which $v^z = 1/24$ which additionally tests the ability of the conservative-to-primitive variable inversion scheme to maintain this quantity unchanged over time. For the three-dimensional case, we set $v^x = 0.2\sqrt{2}$, $v^y = 0.2$, $v^z = 0.1$, and then rotate the tube and all vectors by $\pi/4$ radians in the $x - z$ -plane so that the loop runs along the $x - z$ face diagonal. In all cases, velocity components are chosen to yield integer intervals over which the flux tube should propagate across the grid, $T = 24$ for the two-dimensional cases and

$T = 5$ for the three-dimensional case. Runs were performed using RK2 timestepping, TVD reconstruction with an monotonized central limiter, the HLLC Riemann solver, divergence cleaning active with $\kappa = 10$, and CFL factors of 0.4 for the 2-d cases and 0.25 for the 3-d case.

In the top panels of figure 10, we show the initial state of the flux tube of the two-dimensional case, with B^x on the left and the magnetic pressure $P_m = b^2/2$ on the right. In the centre panels, we show the results of a low-resolution simulation, with $\Delta x = 1/128$, and in the bottom panels a high-resolution simulation with $\Delta x = 1/256$. Both simulations are evolved for one grid crossing time, $T = 24$, with $v^z = 1/24$. We find no substantial differences in the results after one period between $v^z = 0$ and $v^z = 1/24$. While the observable decrease in the central magnetic pressure is the inevitable result of the simulation, we find the effect is significantly decreased by the increase in resolution. We find similar results for the three-dimensional loop, whose evolution is depicted by figure 11, and demonstrates comparable uniformity in the pressure and field strength for a given resolution. From top to bottom, we show the initial state and the states after one and two periods ($T = 5$ and $T = 10$, respectively). The initial pressure, which should remain invariant, is maintained throughout the body of the loop, but decreases in the central region and near the edges, as expected (this result can be compared with figure 4 of [106], with which we find excellent agreement).

5.7. Bondi Accretion

Bondi accretion is a steady-state solution for spherically-symmetric, adiabatic accretion onto a black hole. The Einstein Toolkit implementation of the general relativistic analogue of Bondi's original solution [109] follows the methodology laid out in [110], assuming a fixed background spacetime of a Schwarzschild black hole of mass M_{BH} . The solution is determined by the continuity equation

$$\dot{M} = 4\pi\hat{\rho}\hat{u}\hat{r}^2, \quad (102)$$

and the integrated Euler equation for a polytropic equation of state,

$$\dot{Q} = \hat{h}^2\left(1 - \frac{2M_{\text{BH}}}{\hat{r}} + \hat{u}^2\right), \quad (103)$$

where $\hat{\cdot}$ denotes a variable in Schwarzschild coordinates, h is the specific enthalpy, $u = -u_i r^i$, and c_s the speed of sound. Adding a *radial* magnetic field to the system does not alter the Bondi solution. Following [14], we set the radial magnetic field by prescribing a primitive magnetic field vector of

$$B^i = \frac{B_0 M_{\text{BH}}^2 x^i}{\sqrt{\gamma} r^3}, \quad (104)$$

where $r = \sqrt{x^2 + y^2 + z^2}$ is the coordinate radius and B_0 controls the magnetic field strength. A magnetic field vector of this form by construction satisfies the divergence free constraint (19).

A particular solution is chosen by specifying the BH mass M_{BH} , the mass accretion rate \dot{M} , the location of the sonic point in Schwarzschild coordinates (i.e., the areal




Figure 10. Results of the two-dimensional advected loop test with $v^z = 1/24$. In the top panels, we show the initial state of B^x in greyscale, along with the (uniform) magnetic pressure. In the centre and bottom rows, we show the result of low ($\Delta x = 1/128$) and high ($\Delta x = 1/256$) resolution simulations after one period has passed ($T = 24$). The loss of magnetic pressure in the centre of the loop is greatly reduced in the higher-resolution simulation.




Figure 11. Results of the three-dimensional advected loop test with $v^z \neq 0$. In the top panels, we show the 2d $z = 0$ slices through the initial state of B^x in greyscale, along with the (uniform) magnetic pressure. In the centre and bottom rows, we show the results after one ($T = 5$) and two ($T = 10$) periods, respectively. The flux tube formed by the moving flux loop is oriented diagonally in the domain, hence the slice is elliptical rather than circular. The grid spacing used is $\Delta x = 1/128$. Our result is extremely similar in how well it preserves the magnetic flux and pressure of the loop to the results shown in figure 4 of [106].

radius) \hat{r}_s , the magnetic field strength B_0 and the adiabatic index Γ of the equation of state. Table 2 lists the values of the parameters used in our simulations. We explore various values for B_0 and \dot{M} but keep the sonic radius $\hat{r} = 8 M_{\text{BH}}$ and adiabatic index $\Gamma = 4/3$ fixed.

By requiring the solution to remain smooth at the sonic point, the sound speed and thus velocity of the fluid at that point must satisfy $\hat{u}_s = c_s = M_{\text{BH}}/(2\hat{r}_s)$. From this equation we find the polytropic constant K for the EOS and subsequently from the continuity equation the conserved quantity \dot{Q} . The solution in the remaining domain is then determined by integrating outwards (or inwards) from the sonic point before converting to the desired coordinate system of the evolution. We choose to work in ingoing Kerr-Schild coordinates which are horizon penetrating and therefore allow us to model the flow of matter across the event horizon.

To avoid numerical issues due to the singularity inside of the event horizon we use a transition function

$$\chi(r) = \begin{cases} \frac{1}{2} \left[1 + \tanh \left(\tan \left(\frac{\pi r}{M_{\text{BH}}} - \frac{\pi}{2} \right) \right) \right] & r < M_{\text{BH}} \\ 1 & \text{otherwise} \end{cases} \quad (105)$$

to damp magnetic field $B^i \rightarrow \chi \vec{B}^i$ and velocity $v^i \rightarrow \chi v^i$ to zero near the singularity. Similarly we reduce the mass density, pressure and internal energy density to values corresponding to our atmosphere level ρ_{atmo} which is 10^{-8} of the density at the event horizon.

Our numerical setup follows [43], i.e. we simulate a rectangular box of size $[0, 11] \times [0, 11] \times [-11, 11]$ using resolution of 0.11 for the low resolution and 0.073 for the high resolution run. Since the system is spherical, we used a 90 degree rotational symmetry around the z axis to reduce the computational costs. We use a HLLE Riemann solver, TVD reconstruction using a minmod limiter. We integrate in time using a third-order Runge-Kutta (RK3) integrator with Courant factor 0.25. For this test we employed our constrained transport scheme which we found to be more stable in the strongly magnetically dominated regimes explored by us in this test. After each RK substep we re-set the grid inside of $\hat{r} < 1 < M_{\text{BH}}$ and at the outer boundary to the analytic initial solution, providing Dirichlet type boundary data.

We note that due to our choices of numerical timestepping method, stencil size and resolution grid points outside of the event horizon $\hat{r} = 2 M_{\text{BH}}$ are not affected by grid points inside of $\hat{r} = 1 M_{\text{BH}}$. Thus our method is effectively excising the interior of $r < 1 M_{\text{BH}}$ from the simulated region.

To verify the accuracy of our code, we evolve the analytical initial data for $T = 100 M_{\text{BH}}$ and compare the data along the $z = 0.055$, $y = 0$ line to the analytic initial data. We display the result of this comparison in figure 12. We observe 2nd order convergence with increasing resolution in the region outside of the event horizon at $r = 2 M_{\text{BH}}$. Our results compare very well with those of [43], figure 6.

To adequately test the stability of the steady-state solution, we repeat this test and vary both the accretion rate and the magnetic field strength. We also include a

Table 2. List of tests performed for the magnetised Bondi infall. The first column shows the run name, the second the magnetic field strength parameter, the third the (unitless) mass accretion rate, the fourth the ratio of magnetic pressure to gas density at the horizon, the fifth the resolution and the last column lists the maximum relative error in ρ outside of $r = 2M$ for each test run.

Run	B_0	\dot{M}	b^2/ρ	Δx	$ (\rho/\rho_{\text{exact}}) - 1 _\infty$
A	0.0	12.57	0.0	0.11	$9.70 \cdot 10^{-5}$
B	0.0	12.57	0.0	0.073	$4.60 \cdot 10^{-5}$
C	11.39	12.57	25.0	0.11	$7.42 \cdot 10^{-2}$
D	11.39	12.57	25.0	0.073	$3.92 \cdot 10^{-2}$
E	11.39	25.13	12.5	0.11	$3.88 \cdot 10^{-2}$
F	5.7	12.57	6.26	0.11	$2.21 \cdot 10^{-2}$
G	5.7	25.13	3.13	0.11	$1.37 \cdot 10^{-2}$

pure hydrodynamic simulation for each accretion rate. In Table 2, we list the set of parameters and their values that we vary for this test along the supremum norm of the relative error on the grid outside of the event horizon after $t = 100 M_{\text{BH}}$ of evolution.

We found that higher order reconstruction method such as PPM and methods employing divergence cleaning performed less well in this test. In the case of divergence cleaning we suspect that our inability to specify initial data with vanishing *numerical* divergence is partially the source of these difficulties. Even in the constrained transport case, we observe a certain amount of initial drift away from the analytic solution in the field values. Eventually the solution settles down and the amount drift to decrease with increasing resolution. In the case of divergence cleaning, our imperfect boundary data serves as a constant source for constraint violations, possibly spoiling the result. Similarly the interpolation inherent in adaptive mesh refinement techniques seems to serve as a source of constraint violation and prevent us from simulating this system using AMR and either divergence cleaning or constraint transport.

5.8. Magnetised TOV Star

In this section we consider the stability of a stationary, poloidally magnetised spherically symmetric neutron star model, obtained via the solution of the Tolman-Oppenheimer-Volkhoff equation [111]. We follow [17] and set up a poloidal magnetic field on top of the independently solved fluid configuration. This setup provides a more challenging test for the code than the previous ones since it probes both the GR and the MHD evolution. While the initial configuration is stationary, we evolve both the spacetime and the GRMHD equations to test the ability of the new code to maintain such a solution. The hydrodynamic initial data is a TOV star generated using a polytropic EOS with polytropic constant $K = 100$, adiabatic index $\Gamma = 2.0$ and initial central density of $1.28 \cdot 10^{-3} M_\odot$. The fluid is evolved with a Γ -law EOS and $\Gamma = 2$. As in [12], we introduce a poloidal field that we add on top of the TOV initial data. We choose the vector potential $A_\varphi = A_b \varpi^2 (1 - \rho_0/\rho_0^{\text{max}})^{n_p} \max(P - P_{\text{cut}}, 0)$ where




Figure 12. Bondi inflow after $100M_{\text{BH}}$ of evolution compared with the semi-analytic solution for varying numerical resolution. We observe 2nd order convergence in the region outside of the event horizon at $r = 2M_{\text{BH}}$. The red dashed line shows the errors in ρ compared to the analytical solution for the low resolution run $r0$ with resolution 0.11 while the black solid line shows the error in the high resolution run $r1$ using a resolution of 0.073 scaled so that for second order convergence the curves coincide with each other. This inset shows a zoom in of the region around the event horizon where matter is moving fastest and the magnetic field is strongest, i.e. where we expect the largest errors. Our errors are almost identical to those found by [43].

$\varpi \equiv \sqrt{(x - x_*)^2 + y^2}$. Here A_b determines the strength of the initial magnetic field, n_p the offset of the maximum of the magnetic field strength in respect to the maximum density ρ_{max} of the fluid, and P_{cut} the pressure outside of which the magnetic field is set to zero. In this test we confine the magnetic field to the interior of the star by setting $P_{\text{cut}} = 10^{12}P_{\text{atmo}}$, much higher than the atmospheric pressure. The parameter A_b is set to unity, which, together with $n_p = 0$, provides a magnetic field with a maximum strength of 8.5×10^{-6} in geometrised units which corresponds to 2.5×10^{15} G.

We place the TOV star with radius $8.13M_{\odot}$ and gravitational mass $1.40M_{\odot}$ at the origin of a grid with 5 levels of mesh refinement with cubical extents $320M_{\odot}$, $120M_{\odot}$, $60M_{\odot}$, $30M_{\odot}$, and $15M_{\odot}$. The extent of the finest level is chosen in order to cover the entire star. We perform three simulations with coarsest resolutions of $8M_{\odot}$ ($r0$), $4M_{\odot}$ ($r1$), and $2M_{\odot}$ ($r2$). The corresponding resolutions of the finest grids are $1.0M_{\odot}$ ($r0$), $0.5M_{\odot}$ ($r1$) and $0.25M_{\odot}$ ($r2$). We simulate for 6.9 ms using oPPM reconstruction, the HLLE Riemann solver, and evolve the magnetic field with constrained transport. We use 4th-order Runge-Kutta time integration with a Courant factor of 0.25. We set the Γ -Driver gauge condition [59] parameter to $\eta = 0.5$ and add Kreiss-Oliger dissipation to the spacetime variables with a dissipation coefficient of $\epsilon = 0.1$.

Figure 13 shows the fractional change of the maximum rest-mass density ρ for




Figure 13. Results of the magnetised TOV simulation. First panel: Normalised maximum rest-mass density $\rho_{\max}/\rho_{\max,0}$. Second panel: L_2 -norm of the Hamiltonian constraint $\|H\|_2$ for the three resolutions $r0$, $r1$, and $r2$. Third panel: Absolute difference ϵ in normalised maximum rest-mass density between low ($r0$) and medium ($r1$), and medium ($r1$) and high ($r2$) resolutions are shown (centre panel). Fourth panel: L_2 -norm of the normalised divergence of the magnetic field $|\nabla_i B^i / (|B| \cdot \Delta x)|_2$, where Δx is the finest grid spacing, for each of the three resolutions $r0$, $r1$, and $r2$ and $|\vec{B}|_\infty$ is the maximum magnetic field strength on the grid. The differences and the L_2 -norm of the Hamiltonian constraint and the divergence of magnetic field are scaled for first order convergence.

all three resolutions, the differences between low ($r0$) and medium ($r1$), and medium ($r1$) and high ($r2$) resolutions in maximum rest mass density and the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. The differences in rest-mass density and the Hamiltonian constraint are scaled for first order convergence. While the TOV solution is stationary, numerical perturbations excite oscillations in the central density of the star, commonly observed in simulations of TOV stars, which must converge away with increasing resolution. We can analyse the convergence properties of our simulations more carefully by considering the second panel of figure 13, which shows the absolute differences between low ($r0$) and medium ($r1$), and medium ($r1$) and high ($r2$) resolutions in the normalised maximum rest-mass density. The differences are scaled by a factor of 2 for first-order convergence. We observe between first and second order convergence, which is the expected behaviour, since the oPPM scheme reduces to first-order accuracy at the surface of the star. The third panel of figure 13 shows the L_2 -norm of the evolution of the Hamiltonian constraint $\|H\|_2$ for the resolutions $r0$, $r1$ and $r2$. The medium and high resolutions are rescaled for first-order convergence. We again observe between first-order and second-order convergence, consistent with what is expected. The fourth panel of figure 13 shows the L_2 -norm of the normalised divergence of the magnetic field as a function of time for the resolutions $r0$, $r1$ and $r2$. Again, we observe between first-order and second-order convergence. Note that the lines are constant, since constrained transport preserves the initial constraint violation to round-off.

5.9. Rotating Core Collapse

As a final test, we present results for the collapse of a rotating magnetised stellar core with simplified microphysics. We analyse a number of quantities to demonstrate the ability of our code to simulate magnetised stellar collapse in full GR. For this, we consider model A1 of Shibata *et al.* [19], who carried out simulations in axisymmetry. While our code is 3D, we can limit non-axisymmetric dynamics to even ℓ and m (in terms of spherical harmonics) that are multiples of 4 by evolving only in one octant of the 3D domain and enforcing rotational symmetry on the two inner faces of the octant and reflection symmetry in the z -direction. We model the precollapse star as a $\Gamma = 4/3$ -polytrope with a central density of $\rho_c = 10^{10} \text{ g cm}^{-3}$ and a polytropic constant of $K = 4.897 \cdot 10^{14} [\text{cgs}]$. This corresponds to an ultra-relativistic degenerate electron gas at electron fraction $Y_e = 0.5$. We choose the initial rotation profile as $u^t u_\phi = \varphi_d^2 (\Omega_c - \Omega)$, where Ω_c is the central angular velocity, Ω is the angular velocity, and φ_d is a constant describing the degree of differential rotation. Model A1 of [19] is uniformly rotating ($\varphi_d \rightarrow \infty$) and has a ratio of polar to equatorial radius $R_p/R_e = 0.667$. The initial data were generated using the RNS code [112]. During evolution, we employ a hybrid EOS [113, 114], consisting of a two-piece polytropic cold pressure component P_P and a thermal pressure component P_{th} ,

$$P = P_P + P_{th} . \quad (106)$$




Figure 14. First panel: Convergence of the gravitational wave signal A_2 computed via the quadrupole formula as in [19], and gravitational wave signal from [19] for the corresponding model. Second panel: Central rest-mass density ρ_c . Third panel: L_2 -norm of the Hamiltonian constraint $\|H\|_2$. Fourth panel: L_2 -norm of the normalised divergence of the magnetic field $|\nabla_i B^i / (|B| \cdot \Delta x^{-1})|_2$. All four quantities are shown for the three resolutions $r0$, $r1$, and $r2$. The L_2 -norm of the Hamiltonian constraint and the normalised divergence of the magnetic field in panels three and four are scaled for first order convergence. The constraints initially exhibit second-order convergence, which drops to first order after core bounce generates a shock. The divergence of the magnetic field is kept constant by the constrained transport scheme, unless a change in the grid structure occurs. This causes the jumps in the time evolution of $|\nabla_i B^i / (|B| \cdot \Delta x^{-1})|_2$ since the newly created grids are populated using non-constraint-preserving operators.




Figure 15. Convergence of the maximum value of the magnetic field components B^i for the three different resolutions $r0$, $r1$ and $r2$ in the rotating core collapse simulation. Differences between the three resolutions are most obvious for B^x (top panel), but also present for B^y (middle panel) and B^z (bottom panel). This is however to be expected, as differences in the maximum of the magnetic field components may arise from local turbulent behaviour and depend sensitively on resolution. We observe a similar overall evolution of the magnetic field maxima as in [19]. We note that the differences in evolution for B^x_{\max} and B^y_{\max} result from the fact that we choose to plot the maximum value and not the maximum of the absolute value. As our domain is symmetric under rotations by 90 degrees about the z -axis, there is a sign flip between B^x and B^y in any given quadrant of our domain, i.e. $B^x_{\max} = -B^y_{\min}$.

Following [19], in the cold part, the adiabatic index jumps from $\Gamma_1 = 1.3$ to $\Gamma_2 = 2.5$ at $\rho \geq 2 \times 10^{14} \text{ g cm}^{-3}$ to mimic the stiffening of the nuclear EOS at core bounce. The thermal component of the hybrid EOS is a Γ -law with $\Gamma_{\text{th}} = \Gamma_1$. The specific internal energy is given by $\epsilon = \epsilon_P + \epsilon_{\text{th}}$ and consists of a polytropic and a thermal part. We induce collapse by following the prescription outlined in [19], that is we set the internal energy according to $\epsilon = 3K\rho^{1/3}$ and then determine the pressure as $P = 3K(\Gamma_1 - 1)\rho^{4/3}$. This reduces the pressure from its initial value by a factor of $(1 - 3(\Gamma_1 - 1))$ (i.e. 10% for $\Gamma_1 = 1.3$). Following [19], we assume a poloidal initial magnetic field computed from the vector potential

$$A_\phi = A_b \varpi^2 \max \left[\rho^{1/n_b} - \rho_{\text{cut}}^{1/n_b}, 0 \right], \quad (107)$$

Table 3. Initial parameters and properties of model A1.

Central rest-mass density [g cm^{-3}]	ρ_c	1.0×10^{10}
ADM mass [M_\odot]	M_{ADM}	1.503
Baryonic mass [M_\odot]	M_B	1.503
Ratio of rotational kinetic to potential energy	T/W	8.9×10^{-3}
Equatorial radius [km]	R_e	2267
non-dimensional angular momentum parameter	J/M_{ADM}^2	1.235
Ratio of polar to equatorial radius	z/R_e	0.667
Differential rotation parameter	\tilde{A}	∞
Angular velocity at the rotation axis	Ω_c	4.11
Maximum initial magnetic field strength [G]	$ B^i _{\text{max}}$	7.2×10^{12}
Maximum initial ratio of magnetic and fluid pressure	$(P_{\text{mag}}/P)_{\text{max}}$	2.1×10^{-4}
Ratio of magnetic to rotational kinetic energy	$E_{\text{mag}}/T_{\text{rot}}$	1.9×10^{-3}

with $n_b = 1$, $\rho_{\text{cut}} = 10^{-4}\rho_c$ and $\varpi = \sqrt{x^2 + y^2}$. The parameter A_b is chosen so that the maximum initial magnetic field strength is $7.2 \times 10^{12} \text{G}$. For convenience, we summarise the model parameters in Table 3. Our numerical setup initially consists of 4 levels of nested grids each refined by a factor of 2 in resolution. As the star collapses, and the central density increases, we progressively add 5 additional refinement levels, following the approach taken in [115]. The density thresholds are chosen to be $[8.0 \times 10^{10}, 3.2 \times 10^{11}, 1.28 \times 10^{12}, 5.12 \times 10^{12}, 2.05 \times 10^{13}] \text{g cm}^{-3}$. The final grid structure consists of 9 levels of nested grids with cubical extents $[5669.8 \text{ km}, 3061.7 \text{ km}, 2438 \text{ km}, 1559.2 \text{ km}, 283.5 \text{ km}, 212.6 \text{ km}, 144.7 \text{ km}, 59.1 \text{ km}, 23.6 \text{ km}]$. The extent of the finest level is chosen such that it covers the entire compact remnant (the “proto-NS”) after core bounce. The innermost refinement level for our fiducial setup ($r1$) has a linear resolution of 369.1 m. We consider two more resolutions, $r0$ with fine grid resolution 461.4 m and $r2$ with fine grid resolution 295.3 m. This allows us to study the resolution dependence of our results. We use vertex-centred AMR, constrained transport, and TVD reconstruction with the van Leer monotized central limiter. While GRHydro includes higher-order reconstruction schemes, we use TVD in these simulations to match the numerical methods used in [19] as closely as possible. The time integration is performed using RK4 with a Courant factor of 0.2. The damping coefficient of the Γ -driver gauge condition for the spacetime evolution is set to $\eta = 1/2$. We add 5th-order Kreiss-Oliger dissipation to the spacetime variables with a dissipation coefficient $\epsilon = 0.1$. The atmosphere is chosen to be a factor 10^{10} smaller than the central density. Shibata et al. [19] extract the gravitational wave signal from rotating core collapse and bounce via a quadrupole formula and define a wave amplitude

$$A_2 = \ddot{I}_{xx} - \ddot{I}_{zz} \text{ ,} \quad (108)$$

with the definitions of I_{xx} and I_{zz} given in [19]. Since the gravitational wave signal is a good tracer of the overall dynamics of rotating core collapse, we extract A_2 from our simulations for comparison with [19].

In figure 14, we show convergence for the central density ρ_c , the gravitational wave signal A_2 , the L_2 -norm of the Hamiltonian constraint $||H||_2$, and the normalised divergence of the magnetic field $|\nabla_i B^i / (|B| \cdot \Delta x^{-1})|_2$. The Hamiltonian constraint in the third panel of figure 14 and the normalised divergence of the magnetic field in the fourth panel are scaled for first-order convergence. Before core bounce, second order convergence is achieved, but reduced to first order after bounce. Our numerical scheme reduces to first-order accuracy at shocks, contact discontinuities and local maxima, i.e. at the centre of the proto-NS, which dominates the Hamiltonian constraint violation. The gravitational wave signal A_2 , the central rest-mass density ρ_c , the L_2 -norm of the Hamiltonian constraint and the normalised divergence of the magnetic field are shown for the three resolutions $r0$, $r1$, and $r2$. In the top two panels, we also show the gravitational wave signal and central density evolution as extracted and digitised from figures 2 and 17 of [19]. We observe excellent agreement for the gravitational wave amplitude A_2 with the results of [19]. The central density evolution of [19] reaches a peak value of $6.5 \cdot 10^{14} \text{ g cm}^{-3}$ at bounce while we find a maximum bounce value of $6.4 \cdot 10^{14} \text{ g cm}^{-3}$. The post-bounce evolution differs by an approximately constant offset of $\simeq 6\%$, however the features in the time evolution itself are very similar. We attribute the observed deviations in maximum rest-mass density evolution to differences in our numerical setup.

We find between second-order and first-order convergence for the gravitational wave signal A_2 , the central density ρ_c , and the L_2 -norm of the Hamiltonian constraint $||H||_2$, which is within the expected behaviour of the code. The convergence in both central density and maximum rest-mass evolution are less clean than for the Hamiltonian constraint due to their oscillatory nature. The fourth panel of figure 14 shows the evolution of the L_2 -norm of the normalised divergence of the magnetic field $|\nabla_i B^i / (|B| \cdot \Delta x^{-1})|_2$. The constrained transport scheme used in this test maintains the divergence present in the initial data to round-off precision. Constraint violations appear only when additional grids are added and filled by non-constraint-preserving interpolation.

Figure 15 shows the maximum of the magnetic field components $|B^i|_{\text{max}}$ as a function of time for the three resolutions $r0$, $r1$, and $r2$. We observe resolution-dependent differences in the magnetic field evolution in the postbounce phase. This is particularly obvious for B^x , but also visible in B^z . This behaviour is expected, since the evolution of the magnetic field is very sensitive to resolution due to effects of turbulence and instabilities, most importantly the magnetorotational instability (MRI) [116, 117]. A similar behaviour for the resolution dependence of B^i is shown in figure 6 of [19]. Our results for the postbounce magnetic-field evolution differ quantitatively from [19], which is expected due to our different numerical setup. While our setup consists of different layers of mesh refinement with varying resolution throughout our 3D domain, [19] employ a uniform axisymmetric grid with a coordinate singularity at the rotation axis.

6. Conclusions

We have presented the implementation and a set of rigorous tests of ideal general-relativistic magnetohydrodynamics (GRMHD) in the **GRHydro** code, which is part of the Einstein Toolkit. All tests and results presented in this paper can be reproduced with the code and input files available at <http://einsteintoolkit.org>.

GRHydro is the first open-source code that provides for fully self-consistent GRMHD evolutions on dynamical general-relativistic spacetimes. Its development includes many contributions from a wide array of collaborators involved with other relativistic codes, particularly the original version of the general-relativistic hydrodynamics code **Whisky** [32], from which **GRHydro** started, and the publicly available fixed-background **HARM** code [44], from which we have adapted a number of routines and techniques. Combined with other functionality from the Einstein Toolkit, the code may be used to perform GRMHD evolutions of a broad range of relativistic astrophysical systems. We expect these to primarily consist of magnetised isolated and binary neutron stars and collapsing stellar cores, but should the proper initial data be constructed, accretion disks and tori, white dwarfs, and a number of other configurations could be simulated as well.

Our implementation of the equations of GRMHD also features the ability to eliminate spurious divergence in the magnetic field through two different methods: hyperbolic divergence cleaning and constrained transport. While these are clearly sufficient for many problems, one of the responsibilities in implementing a community-based code is to incorporate techniques developed by other groups to solve similar problems. Along these lines, we foresee work in the future to incorporate vector potential treatments in which the B -field is computed from the values of an underlying vector potential A that serves as a fundamental evolved variable. Such a scheme has recently been implemented by Etienne *et al.* [69, 92] and Giacomazzo *et al.* [118]. The former work also included appropriate electromagnetic gauge conditions for evolutions of compact object binaries involving neutron stars and other relativistic systems of interest.

The **GRHydro** code now includes several improved higher-order techniques for reconstruction of states at cell interfaces for use in high-resolution shock capturing, including three that were described in the original Einstein Toolkit code paper [22], TVD, PPM, and ENO, as well as three new methods that have been implemented and released since, ePPM [36, 78], WENO5 [79, 86] and MP5 [80] (see Appendix B).

In the near future, we will extend the new GRMHD version of **GRHydro** to the generalised multipatch infrastructure of [36], which will be released as part of the Einstein Toolkit and will enable simulations on a combination of Cartesian and curvilinear grids. Further work will be directed towards implementing improved approximate Riemann solvers for GRMHD. While the Roe [119] and Marquina [120, 121] solvers implemented in **GRHydro** for purely hydrodynamic evolutions are likely to be too computationally expensive, requiring an eigenmode breakdown of an 8×8 matrix system at every grid cell, the HLLC solver, which recovers the contact wave

dropped by the HLLE approximation, may be an attractive and computationally efficient improvement [74, 122]. Additionally, we will introduce improved techniques for performing conservative-to-primitive variable conversions to handle cases in which the current routines are known to fail, particularly those where the Lorentz factor becomes very large or the fluid becomes dominated by magnetic pressure [123].

In the long term, we expect the GRHydro code and the Einstein Toolkit to evolve into a toolkit not just for numerical relativity, but for the broad computational relativistic astrophysics community. Towards this end, we will include additional physics capabilities in future versions of the Toolkit. These will include better treatments of finite-temperature and composition dependence, nuclear reaction networks, radiation transport of neutrinos and photons, and improvements over the current ideal GRMHD approximation, following the most recent developments in resistive GRMHD [124, 125].

Acknowledgments

We thank Bruno Giacomazzo for use of his exact Riemann solver code, as well as helpful conversations on a number of topics related to the work here. We also thank E. Abdikamalov, M. Campanelli, M. Duez, Z. Etienne, B. Farris, U. Gamma, P. Laguna, L. Lehner, C. Lousto, C. Palenzuela, V. Paschilidis, J. Penner, and Y. Zlochower for useful conversations related to various aspects of the presented work.

The Einstein Toolkit is directly supported by the National Science Foundation in the USA under the grant numbers 0903973 / 0903782 / 0904015 (CIGR) and 1212401 / 1212426 / 1212433 / 1212460 (Einstein Toolkit). Related grants contribute directly and indirectly to work in support of the Einstein Toolkit, including NSF AST-1028087, NSF PHY-1214449, NSF AST-1212170, CAREER NSF PHY-0969855, NSF OCI-0905046, DMS-0820923, NASA 08-ATFP08-0093, an NSERC grant to Erik Schnetter, and Deutsche Forschungsgemeinschaft grant SFB/Transregio 7 “Gravitational-Wave Astronomy”. Christian Ott acknowledges support from the Alfred P. Sloan Foundation. Roland Haas acknowledges support by the Natural Sciences and Engineering Council of Canada. Christian Reisswig acknowledges support by NASA through Einstein Postdoctoral Fellowship grant number PF2-130099 awarded by the Chandra X-ray center, which is operated by the Smithsonian Astrophysical Observatory for NASA under contract NAS8-03060. Tanja Bode acknowledges support from NSF grants 0941417/12058564.

Portions of this research were conducted with computational resources provided by Louisiana State University (allocations `hpc_cactus`, `hpc_numrel` and `hpc_hyrel`), by the Louisiana Optical Network Initiative (allocations `loni_cactus` and `loni_numrel`), by the National Science Foundation through XSEDE resources (allocations TG-PHY060027N, TG-ASC120003, TG-PHY100033, TG-MCA02N014, and TG-PHY120016), the Georgia Tech FoRCE cluster, and the Caltech compute cluster Zwicky (NSF MRI award PHY-0960291).

The multi-dimensional visualizations were generated with the open-source `VisIt`

Table A1. Overview of differences in variable notation between HARM and Valencia.

Quantity	Valencia/ET	HARM
Pressure	P	p
Energy density	$\rho h \equiv \rho(1 + \epsilon) + P$	w
3-velocity for a normal observer	$v^i \equiv \frac{u^i}{\gamma} + \frac{\beta^i}{\alpha}$	v^i
Lorentz factor	W	γ
Projected velocity	$W v^i$	$\tilde{u}^i \equiv j_\mu^i u^\mu = u^i + \frac{W \beta^i}{\alpha}$
Momentum density for a normal observer	$\mathcal{S}_\mu \equiv \alpha T_\mu^0$	Q_μ
Projected momentum density	$\tilde{\mathcal{S}}_\mu \equiv (\delta_\mu^\nu + n_\mu n^\nu) \mathcal{S}_\nu$	\tilde{Q}_μ
Auxiliary Con2prim variable	$Q \equiv \rho h W^2$	W

visualisation package [50, 126]. All other figures were generated with the Python-based matplotlib package [127].

Appendix A. Notation in HARM and Valencia formulations

The conservative to primitive inversion routine found in the Einstein Toolkit was originally developed as part of the HARM GRMHD code [63], and reflects its notation internally. For historical reasons the development of GRMHD in the Einstein Toolkit follows the Valencia formulation of equations [6, 37–39, 73], and this notation was used throughout the first paper summarising the Einstein Toolkit [22]. To convert from one set of notation to the other, we note that the equivalence of various physical quantities in Table A1.

Appendix B. Comparison of reconstruction methods

We briefly compare the new reconstruction methods (section 4.2) that have been implemented in GRHydro. In [36] a detailed comparison between ePPM and oPPM for cell-centred and vertex-centred AMR was performed. These authors found that ePPM is superior to oPPM in all cases they considered, and especially in those using cell-centred AMR. Here, we compare oPPM, ePPM, WENO5, and MP5 reconstruction using the example of an magnetised dynamically evolved TOV star using vertex-centred AMR. We use 7 levels of mesh refinement with cubical extent $[400M_\odot, 200M_\odot, 50M_\odot, 30M_\odot, 19M_\odot, 12M_\odot]$ and employ octant symmetry. The finest level has a resolution of $\Delta x = 0.2M_\odot$ and contains the entire star. We use the HLLE Riemann solver. We evolve spacetime with a Γ -driver gauge parameter of $\eta = 0.5$ and 5th-order Kreiss-Oliger dissipation with $\epsilon = 0.1$ is added to the right-hand side of spacetime variable. We integrate in time with a fourth-order Runge-Kutta integrator and a Courant factor of 0.4. We construct the initial initial data by solving the TOV equations using a polytropic EOS with $\Gamma = 2$ and scale parameter $K = 100 M_\odot$. The




Figure B1. Comparison of oPPM, ePPM, WENO5, and MP5 reconstruction based on the example of a dynamically evolved TOV star. We show the L_2 -norm of the Hamiltonian constraint $\|H\|_2$. MP5 performs best, while oPPM performs worst in this simple test.

central density is set to $\rho_c = 1.28 \times 10^{-3} M_\odot^{-2}$. During the evolution, we use a Γ -law EOS with $\Gamma = 2$.

In figure B1, we show the L_2 -norm of the Hamiltonian constraint $\|H\|_2$ as a function of time. All simulation settings are identical, except for the reconstruction method. We find that MP5 reconstruction performs better than ePPM or WENO5 in this simple test. WENO5 performs worse than ePPM and MP5. The worst results are obtained with oPPM which reduces to first order at the smooth density maximum at the center of the star. In the vertex-centred case, however, the error due to oPPM reconstruction is not as large as in the cell-centred case [36], since in the former, the central density maximum is approximately located on a grid point. With oPPM, the dominant source of error is thus generated at the center of the star. The other reconstruction methods maintain high order at the smooth density maximum. With them, the dominant source of error is generated at the steep density drop at the surface of the star. We emphasise, however, that this test is only one of many possible tests. A more detailed comparison would be necessary for more general conclusions.

Appendix C. Full set of planar MHD shocktubes

The five test cases collected in [96] are widely used to test new MHD codes. We verified that our code passes all tests and compare the results to the exact solution of [43]. The main text 5.2 contains detailed description of our test setup.

In figures C1 and C2 we present results for the moderate and strong blast wave tests




Figure C1. Evolution of the Balsara2 shocktube case, performed with divergence cleaning, with all conventions as in figure 3. The numerical code reproduces the analytical result very well, with only some oscillations visible in the density ρ near the shock front, capturing the sharp shock features at least as well as [96]. A more detailed description of the test setup and parameters can be found in the main text in section 5.2.

of [96] which prove much more challenging to the code due to the Lorentz contraction of the fast moving shock. Finally in figure C3 we present the non-coplanar test problem of [96].




Figure C2. Evolution of the Balsara3 shocktube case, performed with divergence cleaning, with all conventions as in figure 3. This test proved more demanding on our code, in particular the narrow right moving fast shock is not resolved well using 1600 grid points, resulting in a sizeable overshoot near the shock. This feature is common to other codes [96] independent of the specific formulation used. A more detailed description of the test setup and parameters can be found in the main text in section 5.2.




Figure C3. Evolution of the Balsara5 shocktube case, performed with divergence cleaning, with all conventions as in figure 3. A more detailed description of the test setup and parameters can be found in the main text in section 5.2.

Appendix D. Monopole tests modelling high frequency noise

We also include initial data in our tests that models high-frequency numerical noise in the magnetic field. Here, the spatial dependence of B^x is set as in section 5.1 and the other components of the B -field are set to the same amplitude or to zero, with non-zero components multiplied by $(-1)^{N(i,j,k)}$, where $N(i,j,k)$ is a function of the indices i , j , and k of the grid cell in the x , y , and z -directions, to produce an alternating pattern. For a one-dimensional alternating pattern we set $N(i,j,k) = i + j + k$, and only set B^x to be non-zero and alternating. For a two-dimensional test, we use the same function for $N(i,j,k)$ and apply it to B^x to make it alternate, but set B^y to the original Gaussian state and $B^z = 0$. For a three-dimensional test, we set $N(i,j,k) = i + j$, so the alternation is two-dimensional in character, but set $B^x = B^y = B^z$ and allow all three to alternate.

Figure D1 shows the performance of the divergence cleaning algorithm in removing high frequency constraint violating noise such as noise generated during the numerical evolution. Clearly the algorithm is very efficient in damping out the high frequency constraint violations, reducing them much more quickly than the simple Gaussian pulse present in figure 2.




Figure D1. Behaviour of our divergence cleaning scheme, demonstrated by monopole damping and advection for a three-dimensional alternating Gaussian configuration. Conventions are the same as in figure 2. We still see more significant divergences travelling further from their initial location for $\kappa = 1$, but damping is markedly faster for the case $\kappa = 100$ when dealing with high-frequency noise as compared to the lower-frequency modes characterising the Gaussian monopole, seemingly because of cancellation effects from neighbouring points. The choice $\kappa = 10$ still yields the best results for numerically eliminating spurious divergence of the magnetic field.

References

- [1] Shibata M and Uryu K 2000 *Phys. Rev. D* **61** 064001
- [2] Pretorius F 2005 *Phys. Rev. Lett.* **95** 121101
- [3] Baker J G, Centrella J, Choi D I, Koppitz M and van Meter J 2006 *Phys. Rev. Lett.* **96** 111102
- [4] Campanelli M, Lousto C O, Marronetti P and Zlochower Y 2006 *Phys. Rev. Lett.* **96** 111101
- [5] Ott C D, Reisswig C, Schnetter E, O'Connor E, Sperhake U, Löffler F, Diener P, Abdikamalov E, Hawke I and Burrows A 2011 *Phys. Rev. Lett.* **106** 161103
- [6] Font J A 2008 *Living Rev. Relativity* **11** 7. URL (cited on 01-01-2012): <http://www.livingreviews.org/lrr--2008>
- [7] Balbus S A and Hawley J F 1998 *Rev. of Mod. Phys.* **70** 1
- [8] Hawley J F 2009 *Astrophys. Space. Sci.* **320** 107
- [9] Anderson M, Hirschmann E W, Lehner L, Liebling S L, Motl P M, Neilsen D, Palenzuela C and Tohline J E 2008 *Phys. Rev. Lett.* **100** 191101
- [10] Giacomazzo B, Rezzolla L and Baiotti L 2009 *Mon. Not. Roy. Astron. Soc.* **399** L164
- [11] Giacomazzo B, Rezzolla L and Baiotti L 2011 *Phys. Rev. D* **83** 044014
- [12] Liu Y T, Shapiro S L, Etienne Z B and Taniguchi K 2008 *Phys. Rev. D* **78** 024012
- [13] Chawla S, Anderson M, Besselman M, Lehner L, Liebling S L, Motl P M and Neilsen D 2010 *Phys. Rev. Lett.* **105** 111101
- [14] Etienne Z B, Liu Y T, Paschalidis V and Shapiro S L 2012 *Phys. Rev. D* **85** 064029
- [15] Shibata M, Sekiguchi Y and Takahashi R 2007 *Prog. Theor. Phys.* **118** 257
- [16] Shibata M and Sekiguchi Y 2012 *Prog. Theor. Phys.* **127** 535
- [17] Duez M D, Liu Y T, Shapiro S L, Shibata M and Stephens B C 2006 *Phys. Rev. Lett.* **96** 031101
- [18] Lehner L, Palenzuela C, Liebling S L, Thompson C and Hanna C 2012 *Phys. Rev.* **D86** 104035
- [19] Shibata M, Liu Y T, Shapiro S L and Stephens B C 2006 *Phys. Rev. D* **74** 104026
- [20] Giacomazzo B, Baker J G, Miller M C, Reynolds C S and van Meter J R 2012 *Astrophys. J. Lett.* **752** L15
- [21] Farris B D, Gold R, Paschalidis V, Etienne Z B and Shapiro S L 2012 *Phys. Rev. Lett.* **109** 221102
- [22] Löffler F, Faber J, Bentivegna E, Bode T, Diener P, Haas R, Hinder I, Mundim B C, Ott C D, Schnetter E, Allen G, Campanelli M and Laguna P 2012 *Class. Quantum Grav.* **29** 115001
- [23] Einstein Toolkit: Open Software for Relativistic Astrophysics URL <http://einstein toolkit.org/>
- [24] Cactus Computational Toolkit URL <http://www.cactuscode.org/>
- [25] Goodale T, Allen G, Lanfermann G, Massó J, Radke T, Seidel E and Shalf J 2003 The Cactus framework and toolkit: Design and applications *Vector and Parallel Processing – VECPAR'2002, 5th International Conference, Lecture Notes in Computer Science* (Berlin: Springer) URL <http://edoc.mpg.de/3341>
- [26] Allen G, Goodale T, Lanfermann G, Radke T, Rideout D and Thornburg J 2011 *Cactus Users' Guide* URL <http://www.cactuscode.org/Guides/Stable/UsersGuide/UsersGuideStable.pdf>
- [27] Schnetter E, Hawley S H and Hawke I 2004 *Class. Quantum Grav.* **21** 1465
- [28] Schnetter E, Diener P, Dorband E N and Tiglio M 2006 *Class. Quantum Grav.* **23** S553
- [29] Carpet: Adaptive Mesh Refinement for the Cactus Framework URL <http://www.carpetcode.org/>

- [30] Brown J D, Diener P, Sarbach O, Schnetter E and Tiglio M 2009 *Phys. Rev. D* **79** 044023
- [31] Reisswig C, Ott C D, Sperhake U and Schnetter E 2011 *Phys. Rev. D* **83** 064008
- [32] Baiotti L, Hawke I, Montero P J, Löffler F, Rezzolla L, Stergioulas N, Font J A and Seidel E 2005 *Phys. Rev. D* **71** 024035
- [33] Hawke I, Löffler F and Nerozzi A 2005 *Phys. Rev. D* **71** 104006
- [34] Baiotti L, Hawke I, Montero P J and Rezzolla L 2003 *Mem. Soc. Ast. It.* **74** S210
- [35] Whisky – relativistic hydrodynamics and magnetohydrodynamics URL <http://www.whiskycode.org/>
- [36] Reisswig C, Haas R, Ott C D, Abdikamalov E, Moesta P, Pollney D and Schnetter E 2013 *Accepted for publication in Phys. Rev. D.; arXiv:1212.1191*
- [37] Martí J M, Ibáñez J M and Miralles J M 1991 *Phys. Rev. D* **43** 3794
- [38] Banyuls F, Font J A, Ibanez J M, Marti J M and Miralles J A 1997 *Astrophys. J.* **476** 221
- [39] Ibáñez J, Aloy M, Font J, Martí J, Miralles J and Pons J 1999 Riemann solvers in general relativistic hydrodynamics *Proceedings of Godunov Methods, Theory and Applications 1999* ed Toro E (New York: Kluwer Academic/Plenum Publishers) p 485 ISBN 978-0306466014
- [40] Tóth G 2000 *J. Comp. Phys.* **161** 605
- [41] Liebling S L, Lehner L, Neilsen D and Palenzuela C 2010 *Phys. Rev. D* **81** 124023
- [42] Penner A J 2011 *Mon. Not. Roy. Astron. Soc.* **414** 1467
- [43] Giacomazzo B and Rezzolla L 2007 *Class. Quantum Grav.* **24** S235
- [44] Gammie C F, McKinney J C and Toth G 2003 *Astrophys. J.* **589** 444
- [45] Seidel E L, Allen G, Brandt S R, Löffler F and Schnetter E 2010 Simplifying Complex Software Assembly: The Component Retrieval Language and Implementation *Proceedings of the 2010 TeraGrid Conference (Preprint arXiv:1009.1342[cs.PL])*
- [46] Allen G, Goodale T, Löffler F, Rideout D, Schnetter E and Seidel E L 2010 Component Specification in the Cactus Framework: The Cactus Configuration Language *Grid2010: Proceedings of the 11th IEEE/ACM International Conference on Grid Computing (Preprint arXiv:1009.1341[cs.DC])*
- [47] Thomas M and Schnetter E 2010 Simulation factory: Taming application configuration and workflow on high-end resources *Grid Computing (GRID), 2010 11th IEEE/ACM International Conference on* p 369 (*Preprint arXiv:1008.4571[cs.DC]*)
- [48] SimFactory: Herding numerical simulations URL <http://simfactory.org/>
- [49] Korobkin O, Allen G, Brandt S R, Bentivegna E, Diener P, Ge J, Löffler F, Schnetter E and Tao J 2011 Runtime analysis tools for parallel scientific applications *Proceedings of the 2011 TeraGrid Conference: Extreme Digital Discovery TG '11* (New York, NY, USA: ACM) p 22:1:8 ISBN 978-1-4503-0888-5
- [50] Childs H, Brugger E S, Bonnell K S, Meredith J S, Miller M, Whitlock B J and Max N 2005 A contract-based system for large data visualization *Proceedings of IEEE Visualization 2005* p 190
- [51] Berger M J and Olinger J 1984 *J. Comput. Phys.* **53** 484
- [52] Husa S, Hinder I and Lechner C 2006 *Comput. Phys. Commun.* **174** 983
- [53] Lechner C, Alic D and Husa S 2004 *Analele Universitatii de Vest din Timisoara, Seria Matematica-Informatica* **42**
- [54] Kranc: Kranc assembles numerical code URL <http://kranccode.org/>
- [55] Nakamura T, Oohara K and Kojima Y 1987 *Prog. Theor. Phys. Suppl.* **90** 1
- [56] Shibata M and Nakamura T 1995 *Phys. Rev. D* **52** 5428
- [57] Baumgarte T W and Shapiro S L 1999 *Phys. Rev. D* **59** 024007

- [58] Alcubierre M, Brügmann B, Dramlitsch T, Font J A, Papadopoulos P, Seidel E, Stergioulas N and Takahashi R 2000 *Phys. Rev. D* **62** 044034
- [59] Alcubierre M, Brügmann B, Diener P, Koppitz M, Pollney D, Seidel E and Takahashi R 2003 *Phys. Rev. D* **67** 084023
- [60] Hyman J M 1976 The method of lines solution of partial differential equations Tech. Rep. COO-3077-139 ERDA Mathematics and Computing Laboratory, Courant Institute of Mathematical Sciences, New York University
- [61] Harm2d Code URL <http://rainman.astro.illinois.edu/codelib/codes/harm/src/>
- [62] Noble S C, Krolik J H and Hawley J F 2009 *Astrophys. J.* **692** 411
- [63] Noble S C, Gammie C F, McKinney J C and Del Zanna L 2006 *Astrophys. J.* **641** 626
- [64] Stone J M, Gardiner T A, Teuben P, Hawley J F and Simon J B 2008 *Astrophys. J. Suppl.* **178** 137
- [65] Athena code URL <https://trac.princeton.edu/Athena>
- [66] Del Zanna L, Zanotti O, Bucciantini N and Londrillo P 2007 *Astron. Astrophys.* **473** 11
- [67] Shibata M and Sekiguchi Y i 2005 *Phys. Rev. D* **72** 044014
- [68] Duez M D, Liu Y T, Shapiro S L and Stephens B C 2005 *Phys. Rev. D* **72** 024028
- [69] Etienne Z B, Liu Y T and Shapiro S L 2010 *Phys. Rev. D* **82** 084031
- [70] Anderson M, Hirschmann E, Liebling S L and Neilsen D 2006 *Class. Quant. Grav.* **23** 6503
- [71] Arnowitt R L, Deser S and Misner C W 2008 *General Relativity and Gravitation* **40** 1997
- [72] York Jr J W 1983 The initial value problem and dynamics *Gravitational Radiation* ed N Deruelle & T Piran (North-Holland Publishing Co.) pp 175–201
- [73] Anton L, Zanotti O, Miralles J A, Martí J M, Ibanez J M *et al.* 2006 *Astrophys. J.* **637** 296
- [74] Toro E F 1999 *Riemann Solvers and Numerical Methods for Fluid Dynamics* (Berlin: Springer)
- [75] Colella P and Woodward P R 1984 *J. Comp. Phys.* **54** 174
- [76] Harten A, Engquist B, Osher S and Chakravarthy S R 1987 *J. Comp. Phys.* **71** 231
- [77] Shu C W 1999 *High order methods for computational physics* ed Barth T J and Deconinck H A (New York: Springer) pp 439–582
- [78] McCorquodale P and Colella P 2011 *Comm. Appl. Math. Comp. Sci.* **6** 1
- [79] Shu C W 1998 *Lect. Not. Math.* **1697** 325
- [80] Suresh A and Huynh H T 1997 *J. Comp. Phys.* **136** 83
- [81] Harten A, Lax P D and van Leer B 1983 *SIAM review* **25** 35
- [82] Einfeldt B 1988 *SIAM J. Numer. Anal.* **25** 294
- [83] Giacomazzo B 2007 *General Relativistic Magnetohydrodynamics: fundamental aspects and applications* Ph.D. thesis SISSA Scuola Internazionale Superiore di Studi Avanzati Trieste, ITALY URL <http://digitalibrary.sissa.it/handle/1963/1916>
- [84] Roe P L 1986 *Ann. Rev. Fluid Mech.* **18** 337
- [85] van Leer B 1977 *J. Comp. Phys.* **23** 263
- [86] Tchekhovskoy A, McKinney J C and Narayan R 2007 *Mon. Not. Roy. Astron. Soc.* **379** 469
- [87] Harten A 1983 *J. Comp. Phys.* **49** 357
- [88] Mignone A, Tzeferacos P and Bodo G 2010 *J. Comput. Phys.* **229** 5896
- [89] Penner A J 2011 *Numerical Analysis of the Ultrarelativistic and Magnetized Bondi-Hoyle Problem* Ph.D. thesis University of British Columbia Vancouver, BC, Canada URL <http://hdl.handle.net/2429/36734>
- [90] GRHydro documentation URL <http://einstein toolkit.org/documentation/ThornGuide/ThornGuidech152.html#x171-2681000R1>

- [91] Dedner A, Kemm F, Kröner D, Munz C D, Schnitzer T and Wesenberg M 2002 *J. Comp. Phys.* **175** 645–673
- [92] Etienne Z B, Paschalidis V, Liu Y T and Shapiro S L 2012 *Phys. Rev. D* **85** 024013
- [93] Evans C R and Hawley J F 1988 *Astrophys. J.* **332** 659–677
- [94] Yee K 1966 *Antennas and Propagation, IEEE Transactions on* **14** 302
- [95] Balsara D S and Spicer D S 1999 *J. Comp. Phys.* **149** 270
- [96] Balsara D 2001 *Astrophys. J. Suppl.* **132** 83
- [97] Brio M and Wu C C 1988 *J. Comput. Phys.* **75** 400
- [98] Giacomazzo B and Rezzolla L 2006 *J. Fluid Mech.* **562** 223
- [99] Giacomazzo B and Rezzolla L An Exact Riemann Solver for Relativistic MHD URL <http://www.brunogiacomazzo.org/exactriemannsolv.html>
- [100] Balsara D S 1998 *Astrophys. J. Suppl.* **116** 133 URL <http://stacks.iop.org/0067-0049/116/i=1/a=133>
- [101] Del Zanna L, Bucciantini N and Londrillo P 2003 *Astron. Astrophys.* **400** 397
- [102] Komissarov S S 1999 *Mon. Not. R. Astron. Soc.* **303** 343
- [103] Kreiss H and Oliger J 1973 *Methods for the Approximate Solution of Time Dependent Problems (Global Atmospheric Research Programme (GARP): GARP Publication Series vol 10)* (GARP Publication)
- [104] Runge C 1895 *Mathematische Annalen* **46** 167 URL <http://dx.doi.org/10.1007/BF01446807>
- [105] Kutta W 1901 *Z. Math. Phys.* **46** 435
- [106] Beckwith K and Stone J M 2011 *Astrophys. J. Suppl.* **193** 6
- [107] Devore C R 1991 *J. Comput. Phys.* **92** 142
- [108] Tóth G and Odstrcil D 1996 *J. Comput. Phys.* **128** 82–100
- [109] Bondi H 1952 *Mon. Not. Roy. Astron. Soc.* **112** 195
- [110] Shapiro S and Teukolsky S 1983 *Black holes, white dwarfs, and neutron stars: The physics of compact objects* (New York, Chichester Brisbane, Toronto, Singapore: Wiley & Sons)
- [111] Tolman R C 1939 *Phys. Rev.* **55** 364–373
- [112] Stergioulas N and Friedman J L 1995 *Astrophys. J.* **444** 306
- [113] Janka H T, Zwerger T and Mönchmeyer R 1993 *Astron. Astrophys.* **268** 360
- [114] Dimmelmeier H, Font J A and Müller E 2002 *Astron. Astrophys.* **393** 523
- [115] Ott C D, Dimmelmeier H, Marek A, Janka H T, Hawke I, Zink B and Schnetter E 2007 *Phys. Rev. Lett.* **98** 261101
- [116] Balbus S A and Hawley J F 1991 *Astrophys. J.* **376** 214–233
- [117] Balbus S A and Hawley J F 1998 *Rev. of Mod. Phys.* **70** 1–53
- [118] Giacomazzo B, Rezzolla L and Baiotti L 2011 *Phys. Rev. D* **83** 044014
- [119] Roe P L 1981 *J. Comp. Phys.* **43** 357
- [120] Donat R and Marquina A 1996 *J. Comp. Phys.* **125** 42
- [121] Aloy M, Ibanez J, Martí J and Müller E 1999 *Astrophys. J. Suppl.* **122** 151
- [122] Mignone A and Bodo G 2006 *Mon. Not. R. Astron. Soc.* **368** 1040
- [123] Mignone A and McKinney J C 2007 *Mon. Not. R. Astron. Soc.* **378** 1118–1130
- [124] Dionysopoulou K, Alic D, Palenzuela C, Rezzolla L and Giacomazzo B 2012 (*Preprint arXiv:1208.3487[gr-qc]*)
- [125] Palenzuela C 2012 (*Preprint arXiv:1212.0130[astro-ph.HE]*)
- [126] VisIt visualization tool URL <https://wci.llnl.gov/codes/visit/>
- [127] Hunter J D 2007 *Comput. Sci. Eng.* **9** 90